

# OriNet: A Fully Convolutional Network for 3D Human Pose Estimation

Chenxu Luo<sup>1</sup>  
chenxuluo@jhu.edu  
Xiao Chu<sup>2</sup>  
chuxiao@baidu.com  
Alan Yuille<sup>1</sup>  
ayuille1@jhu.edu

<sup>1</sup> Department of Computer Science  
The Johns Hopkins University  
Baltimore, MD 21218, USA  
<sup>2</sup> Baidu Research (USA)  
Sunnyvale, CA 94089, USA

---

## Abstract

In this paper, we propose a fully convolutional network for 3D human pose estimation from monocular images. We use limb orientations as a new way to represent 3D poses and bind the orientation together with the bounding box of each limb region to better associate images and predictions. The 3D orientations are modeled jointly with 2D keypoint detections. Without additional constraints, this simple method can achieve good results on several large-scale benchmarks. Further experiments show that our method can generalize well to novel scenes and is robust to inaccurate bounding boxes.

## 1 Introduction

Estimating 3D human pose from a single RGB image is a fundamental yet challenging problem. It is potentially useful in many real world applications, such as human-robot interaction, augmented reality and character control. Besides the inherent challenges in 2D pose estimation, 3D pose estimation from monocular images is considered more difficult due to the loss of depth and scale ambiguity.

With the advent of deep neural networks, we saw significant progress in monocular 3D human pose estimation. Currently, the end-to-end training method [23, 35] achieve superior results on standard benchmarks [10]. However, there are several limitations that hinder them to real world settings. First, due to the scale and depth ambiguity from a single image, some works [14, 35] even require a fixed image scale, making it less flexible to generalize to other datasets.

Second, currently most methods require a tightly cropped box around the subject. One reason is that most methods use a fully-connected layer to regress the joint coordinates directly, which makes the network sensitive to backgrounds. Another practical issue is the lack of diverse training data as for 2D human pose estimation [0, 13].

Recently, VNect [16] shows promising results of fully convolutional network for 3D human pose estimation. In that work, they attach the 3D joint coordinates to the corresponding locations in the image. However, the performance is not as good as its fully-connected counterparts. One reason may lie on that the network need to localized another joint (the root

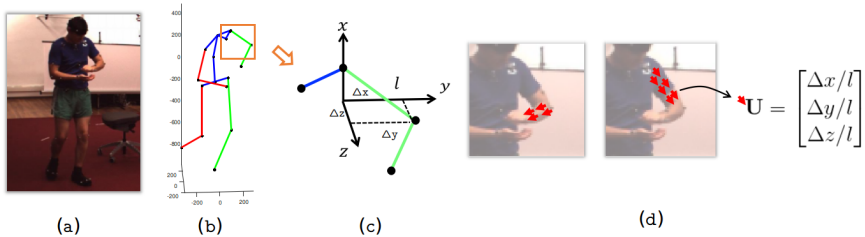


Figure 1: Modeling limb orientations. (a) is an input example. (b) is the 3D human pose represented with skeleton. The key-points are connected following the tree structure, (c) zooming into to see the orientation of the upper arm.  $\Delta x, \Delta y$  and  $\Delta z$  are the relative coordinate of the two keypoints,  $l$  is the length of the limb. (d) showing two examples where orientation vectors are binding with the segmentation of the limb.

joint) when predicting the coordinates of each joint. Given that the actual receptive fields of an FCN are much less than the fully connected networks. This makes it harder to regress 3D coordinates of joints far away from the root joint.

In this paper, we propose a novel fully convolutional network to address the above issues. Instead of directly dealing with joint coordinates, we model limb orientations as a new representation for 3D pose reconstruction. One advantage is that orientation is scale invariant and independent of dataset, which helps resolve the scale ambiguity issue and generalize to diverse data. Since the orientations of limbs are what truly differentiate one pose from another, it is natural and explainable to model the orientations of limbs. Also, it allows the network to focus on pose itself. This representation is flexible and is especially useful for applications such as character control and motion retargeting. Another motivation for using only orientation is that the limb length ratios of different subjects are very similar and are often used as an regularization [27, 28]. So decoupling them to the post-processing stage can let the network focus on each limb without the need to consider other limbs that may lie far away. In our experiments, we show that this is one of the key components to achieve good performance and generalize well to other datasets.

Instead of using a vector representation that lacks spatial association between images and predictions, we propose to combine the orientation with the approximate bounding box of limbs, as shown in Fig.1 (d). Unlike semantic part segmentation [10, 29], we do not care about the detailed boundary of each body part, so a bounding box is a good enough approximate representation for each limb. On the orientation map, regions corresponding to limbs are filled with an orientation vectors, while other locations are set to zero to indicate background. The bounding box of each limb can be easily obtained from 2D joints annotations. This representation can preserve the spatial layout of each limb and explicitly tells the network where to focus when predicting each limb orientation.

Generally, the proposed method can be plugged into any networks for 2D pose estimation. In this paper, we adopt the Stacked Hourglass network [18] because of its superior performance. We simply consider the orientations between adjacent joints without incorporating additional prior knowledge, we show that our method can still achieve the comparable results on Human 3.6m dataset [10] and start-of-the-art results on MPI-INF-3DHP dataset [15]. We expect further constraints such as bone structures [9, 15, 23] or joint-angle limits [10] can improve the performance, but they are beyond the scope of this paper.

Our contributions can be summarized as follow: (1) We propose a fully convolutional network for 3D human pose estimation, which is less sensitive to backgrounds and inaccurate bounding boxes. (2) We propose to use 3D orientations of limbs as a new way to represent 3D poses, which is arguably natural and interpretable. (3) Our proposed method, by only using limb orientations, achieves state-of-the-art results and can generalize well to novel scenes.

## 2 Related Work

Given the difficulties of estimating 3D pose from a single image, many works decouple the problem into two steps: first estimate 2D joints and then lift them into 3D. A typical approach uses sparse-based representation [20, 63]. Recent works also use deep network to regress 3D poses from 2D positions directly [12]. Thanks to the deep networks, we see huge improvements on 2D pose estimation in recent years [3, 5, 6, 7, 18, 26, 28]. The progress also benefits such two-stage methods [33]. However, these two-step methods rely highly on the results of 2D estimation. Also, discarding image cues makes the problem ill-posed.

Recently, some works try to combine the two steps together. In [32], they treat the 2D estimation as latent variables and use the EM algorithm to update the 2D joints at the same time. Tome *et al.* [25] refine both the 2D and the 3D estimation iteratively at each stage. However, 3D poses are still estimated merely from the intermediate 2D results.

Recent works [23, 24, 65] combine the 2D heatmaps and image cues. This method also needs a lot of training data. Since most of the current datasets only contain indoor scene with limited number of subject and background, some methods [23, 65] mix 2D and 3D data during the training time. This makes it less flexible and often need manually correct the discrepancy [65] between dataset annotations.

Most of the above methods use 3D coordinates to represent 3D poses. The coordinates can be relative positions to a root joint [12, 16, 25], to the adjacent joint [12] or their combinations [15]. Another way is to use 2D pixel coordinates and depth to represent each joint [19, 23] and assuming known intrinsic parameter to recover the final 3D pose.

A few works model limbs for pose estimation. In [23] they define bone errors along the skeleton chain. Different from theirs, our method only consider each limb independently without long-range dependencies, making it more flexible to deal with rare poses. Zhou *et al.* [34] also tries to deal with joint angles. They define a set of joint angles as intermediate representation and use a kinematic layer to reconstruct 3D poses. Here we only use the orientations as output. We show that this simple method can achieve better results.

To best of our knowledge, only two kinds of works use fully convolutional networks for 3D pose estimation. Pavlakos *et al.* [19] introduces 3D heatmap to predict per-voxel likelihood by discretizing depth values. But they need the groundtruth depth of a root joint to finally recover the 3D pose, making it less practical. VNect [16] attach the 3D coordinates to a neighbor of the joint on the heatmap. Both of them only deal with each joint separately, without considering the limb orientation which is the underlying factor that causes different poses. Given the flexibility of FCNs, in this paper we propose a more powerful methods for 3D pose estimation.

The orientation has only been explored in the context of 2D pose [3] previously. However, it only serves as an auxiliary task to differentiate different instances, so there is no quantitative evaluation for that. Also it is nontrivial whether it can be extended to estimating 3D orientations accurately, for 3D poses need reasoning beyond the image space. The

motivation and application are totally different from ours.

### 3 Method

In this section, we introduce our fully convolutional network for 3D human pose estimation in detail. Our method model the 2D key-point positions together with the orientation of limbs. The orientation is formulated with an approximate bounding box of the corresponding limb region. An overview of our framework is shown in Fig. 2.

#### 3.1 Orientation Representation

The orientation of each limb is the most discriminative property of 3D poses. It can be represented by a unit vector  $\mathbf{U} = (\Delta x, \Delta y, \Delta z)/l$ , where  $\Delta x, \Delta y$  and  $\Delta z$  are the relative positions of the two joints.  $l = \sqrt{\Delta x^2 + \Delta y^2 + \Delta z^2}$  is the length of the limb. The normalized orientations remove the influence of different human scales and image resolutions.

In order to preserve this spatial layout and explicitly tell the network where to focus, we propose to model the orientation together with the limbs region, as shown in Fig. 2(e). For each limb  $k$ , we generate a bounding box using the locations of two endpoint joints correspond to that limb on the label map. Specifically, the bounding box region  $\mathcal{L}_k$  contains pixels within a predefined width  $w_k$  to the line segment  $\mathbf{p}_{k_1}\mathbf{p}_{k_2}$  of two joints  $k_1$  and  $k_2$ . This forms a good approximation for each limb region, see Fig. 2(e) and supplementary material for detail. Pixels inside the bounding box indicate the region of this limb, and are labeled with the orientation vectors. Other places are set to zero, indicating background. The orientation map is defined as:

$$\widehat{\mathbf{O}}_k(i, j) = \begin{cases} \mathbf{U}_k & (i, j) \in \mathcal{L}_k, \\ \mathbf{0} & (i, j) \notin \mathcal{L}_k, \end{cases} \quad (1)$$

where  $(i, j)$  are the pixel locations on the output prediction map,  $\mathcal{L}_k$  are the regions corresponding to the limb  $k$ . Examples of feature map are shown in Fig. 2.

The loss function for limb orientations is as follows,

$$\mathbf{L}_o = \sum_k \|\mathbf{O}_k - \widehat{\mathbf{O}}_k\|_2^2 \quad (2)$$

where  $\mathbf{O}_k$  is the predicted orientation map for the  $k$ -th limb.

#### 3.2 Model the 2D keypoint location

In order to encourage the connection between the 2D appearance and 3D orientations, we propose to predict the 2D keypoint at the same time. The 2D joints can help localize the limb region more precisely during both training and inference.

Following the common practice, we represent 2D keypoints with heatmaps, as shown in Fig. 2. In each heatmap, a Gaussian centered at ground-truth location indicate the existence of that keypoint. We train 2D keypoints with sigmoid cross entropy loss:

$$\mathbf{L}_p = -\frac{1}{N} \sum_n \sum_{i=1}^W \sum_{j=1}^H \mathbf{p}_{i,j}^n \log(\widehat{\mathbf{p}}_{i,j}^n) + (1 - \mathbf{p}_{i,j}^n) \log(1 - \widehat{\mathbf{p}}_{i,j}^n) \quad (3)$$

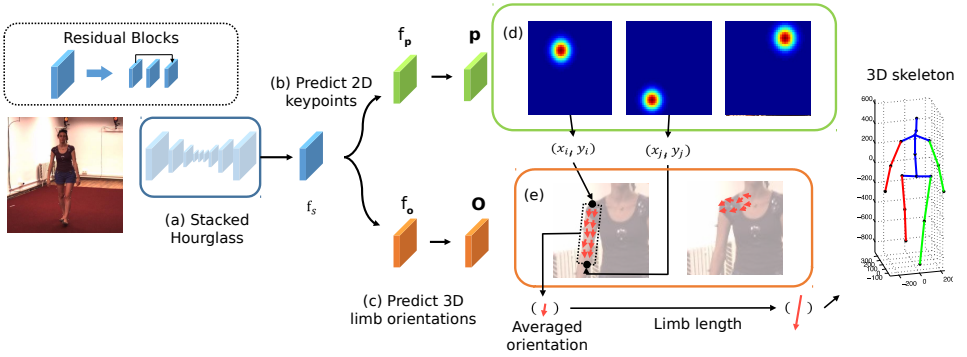


Figure 2: Overview of our framework. We use both 2D heatmaps (d) and limb orientations (e) as the supervisions. The 3D orientation maps are derived from keypoints without extra annotations. In the test stage, the predicted keypoint locations are used to crop the limb regions on the orientation heatmap to get the orientation prediction.

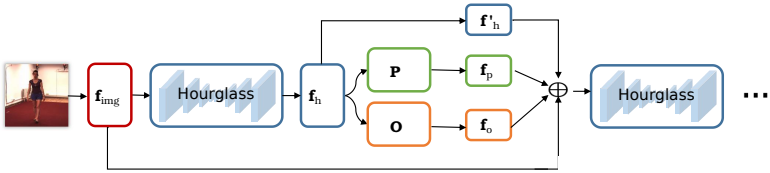


Figure 3: Network structure.  $f_h$  is the output of the hourglass module,  $P$  is the 2D heatmaps and  $O$  is orientation maps.  $f_{img}$  is the image feature and  $f_p$ ,  $f_o$  and  $f'_h$  are the intermediate features,

where  $p^n$  and  $\hat{p}^n$  are predicted and groundtruth heatmaps and  $N$  is the number of joints.

The final loss function is

$$L = L_o + \lambda L_p \quad (4)$$

where  $\lambda$  is a balancing parameter. We use  $\lambda = 0.2$  in the experiments.

### 3.3 Network Structure

We implemented the model based on the Stacked Hourglass network [18]. The network learns low level image representations  $f_{img}$  for the input image. In each stack, the hourglass module will learn its feature  $f_h$ . After that, the network is split into two branches, one for 2D joint locations, the other for the 3D limb orientations (see Fig. 2). Supervision is applied after each stack.

The predictions, hourglass features and the image features are combined as follows:

$$f'_{img} = f_{img} + \text{Conv}(p) + \text{Conv}(O) + \text{Conv}(f_h) \quad (5)$$

Then the updated image feature  $f'_{img}$  is used as the input of next hourglass module. For simplicity, we draw the Hourglass network structure for 2 stacks in Fig. 3.

### 3.4 Inference

The inference process includes (1) estimating 2D keypoints, (2) using the estimated 2D keypoints to read off limb regions, (3) taking the averaged orientations in that region, (4) recovering the 3D pose using the estimated orientations, limb length ratio and scale information.

The 2D keypoint locations can be obtained by finding the maximal location on the 2D key-point heatmap, for the  $n$ -th joint,

$$(x_n, y_n) = \arg \max \mathbf{p}^n \quad (6)$$

After getting predictions of all 2D keypoints, we can define the region for each limb. Taking a pair of adjacent joints  $\mathbf{p}_1$  and  $\mathbf{p}_2$  that corresponds to limb  $l$ , we use the coordinates to crop out the region between them on the limb 3D orientation  $\mathbf{O}_l$ . The cropped regions contain the predictions of the orientation of the  $l$ -th limb.

We average the value of the normalized predictions and then normalize it as the estimated orientation. Then the limb length ratio as used in many other works [24, 65] and scale information [14, 23, 32, 65] are used to recover each limb vector. By choosing a root joint, we reconstruct each joint position along the tree structure iteratively. Indeed this simple representation could cause error accumulation along the path, however, we show that this can still achieve good results without adding additional constraints.

## 4 Experiments

### 4.1 Datasets and Implementation Detail

We evaluate our method on two datasets: Human3.6m [10] and MPI-INF-3DHP [15].

**Human3.6M.** This is currently the largest 3D pose dataset, containing 11 subjects performing 15 actions. We adopt a commonly used protocol: five subjects(1,5,6,7,8) for training and two subjects (9,11) for testing. The videos are down sampled to 10fps. Mean per joint position error(MPJPE) is computed with the root joint aligned. We also report the results after Procrustes alignment, which only focus on the structure of the 3D poses. Following [2], we test our algorithm on all 17 joints defined in [10].

**MPI-INF-3DHP.** This is a newly released dataset. It is more challenging, for it contains more diverse motions and aims for testing the generalization of various methods. Here, we only use the provided training data and do **not** use any background augmentation like [15, 16]. We report the PCK with a threshold of 150mm and AUC on all 2929 testing images.

**Implementation Detail.** We adopt a 5-stack hourglass network [18] and implement it in Torch7 [8]. The initial learning rate is  $2.5e-4$  and decreases by a factor of 10 after every two epochs. We train the model for about 8 epochs, using RMSprop optimizer. We also add scale and color jittering as data augmentation. During the testing time, for fair comparison, we assume known bounding box of the person as in all the other works. We only use a single crop, without using flipping or multiple scale fusion. For the scale issue, we use training subjects to compute the limb-length ratio and scale. We also test our model directly with randomly shifted bounding box randomly as in [16]. Our approach can run at 20fps on a Titan XP.

## 4.2 Results on Human3.6M Dataset

Results are shown in Table 1. For model trained only on Human3.6m dataset, we achieve the state-of-the-art result. Especially for difficult poses like sitting down, which have fewer examples in the training set, our method outperforms other methods by a large margin. This shows that our method is more data efficient. As for models pretrained or mixed trained with 2D datasets (e.g. MPII), our method can still achieve comparable results. Compared with [14, 23], our method has lower error after alignments while slightly higher errors before alignment. Note that [23] also use 2D data during training and camera parameters during testing. The result shows that our predicted poses have the most similar structure compared with the groundtruth. We also report results using the groundtruth limb length, similar to the skeleton-fitting step used in VNect [16]. In Fig. 4, we show some qualitative results

	Direct	Diss.	Eat	Greet	Phone	Photo	Pose	Purch.	Sit	SitD.	Smoke	Wait	WalkD	Walk	WalkT	Ave
Train on Human3.6M from scratch(or pretrained on ImageNet [22])																
[23]	90.2	95.5	82.3	85.0	87.1	94.5	87.9	93.4	100.3	135.4	91.4	87.3	90.4	78.0	86.5	92.4
[14]	<b>64.9</b>	<b>73.5</b>	76.8	86.4	86.3	110.7	<b>68.9</b>	<b>74.8</b>	110.2	173.9	84.9	85.8	86.3	71.4	73.1	88.4
Ours	68.4	77.3	<b>70.2</b>	<b>71.4</b>	<b>75.1</b>	<b>86.5</b>	69.0	76.7	<b>88.2</b>	<b>103.4</b>	<b>73.8</b>	<b>72.1</b>	<b>83.9</b>	<b>58.1</b>	<b>65.4</b>	<b>76.0</b>
Pretrained on 2D pose datasets(e.g. MPII)																
[14]	59.7	69.7	68.8	68.8	76.4	85.2	59.0	75.0	96.2	122.9	70.8	68.5	82.0	54.4	59.8	74.1
[24]	54.2	61.4	60.2	61.2	79.4	78.3	63.1	81.6	<b>70.1</b>	107.3	69.3	70.3	74.3	51.8	63.2	69.7
[14]	<b>51.8</b>	<b>56.2</b>	58.1	<b>59.0</b>	69.5	78.4	<b>55.2</b>	<b>58.1</b>	74.0	<b>94.6</b>	<b>62.3</b>	59.1	<b>65.1</b>	49.5	52.4	<b>62.9</b>
Ours	53.5	60.9	<b>56.3</b>	59.1	<b>64.3</b>	<b>74.4</b>	55.4	63.4	74.8	98.0	61.1	<b>58.2</b>	70.6	<b>49.1</b>	55.7	63.7
Ours*	49.2	57.5	53.9	55.4	62.2	73.9	52.1	60.9	73.8	96.5	60.4	55.6	69.5	46.6	52.4	61.3
Results after Procrustes alignment																
[14]	69.5	80.2	78.2	87.0	100.7	102.7	76.0	69.6	104.7	113.9	89.7	98.5	79.2	82.4	77.2	87.3
[23]	42.1	44.3	45.0	<b>45.4</b>	51.5	<b>53.0</b>	43.2	41.3	59.3	73.3	51.0	44.0	48.0	38.3	44.8	48.3
[14]	<b>39.5</b>	<b>43.2</b>	46.4	47.0	51.0	56.0	<b>41.4</b>	<b>40.6</b>	56.5	<b>69.4</b>	49.2	45.0	49.5	38.0	43.1	47.7
Ours	40.8	44.6	<b>42.1</b>	45.1	<b>48.3</b>	54.6	41.2	42.9	<b>55.5</b>	69.9	<b>46.7</b>	<b>42.5</b>	<b>48.0</b>	<b>36.0</b>	<b>41.4</b>	<b>46.6</b>

Table 1: Mean per joint position errors (MPJPE) in mm on Human3.6M. \* denotes using the groundtruth length for each limb.

from Human3.6M dataset. For each example, we illustrate the image, our result and the groundtruth respectively.

## 4.3 Results on MPI-INF-3DHP

### 4.3.1 Cross Datasets Test

Performance on Human3.6M seems to be saturating. However, the generalization problem largely remain unexplored. MPI-INF-3DHP provides a good testbed for such purpose. We apply models trained on Human3.6m directly to this dataset and results are shown in table 2.

Our method can achieve comparable results with [15], which is specifically designed for transfer learning. Even compared with mixed-training method [65], we still get comparable results. However, as pointed out in [65], they need to manually fix the position of some joints in an ad-hoc manner because of the discrepancy between different annotations.

We also test the direct regression method [14] using 2D groundtruth as input and rescale the output to the universal skeleton. We see that it can not generalize well to different camera viewpoints. Our method can deal with different camera viewpoints and background better.

Method	Training Data	PCK				AUC
		GS	noGS	Outdoor	ALL	ALL
Meta [15]	H36m	70.8	62.3	58.5	64.7	31.7
Zhou [65]	H36m	45.6	45.1	14.4	37.7	20.9
Martinez* [14]	H36m	49.8	42.5	31.2	42.5	17.0
Ours-len*	H36m	65.7	56.0	60.3	60.9	29.9
Our	H36m	69.8	58.3	65.5	64.6	32.1
Ours*	H36m	<b>71.3</b>	59.4	<u>65.7</u>	<u>65.6</u>	<b>33.2</b>
Zhou [65]	H36m+MPII	<u>71.1</u>	<b>64.7</b>	<b>72.7</b>	<b>69.2</b>	<u>32.5</u>

Table 2: Results on MPI-INF-3DHP test set by scene. All models are pretrained on MPII. Our method can achieve comparable results even with mixed training method [65].

	data	Walk	Exe.	Sit	Reach	Floor	Sport	Misc	Total		
		PCK	PCK	PCK	PCK	PCK	PCK	PCK	PCK	AUC	MPJPE
Mehta [15]	(MPII) 3DHP <sup>a</sup>	86.6	75.3	74.8	73.7	52.2	82.1	77.5	75.7	39.3	117.6
	+H3.6M	-	-	-	-	-	-	-	76.5	40.8	-
Mehta [14]	(MPII+LSP) H3.6M+3DHP <sup>a</sup>	87.7	77.4	74.7	72.9	51.3	83.3	80.1	76.6	40.4	124.7
Ours	(MPII)	<b>90.4</b>	<b>79.1</b>	<b>88.5</b>	<b>81.6</b>	<b>66.3</b>	<b>91.9</b>	<b>92.2</b>	<b>81.8</b>	<b>45.2</b>	<b>89.4</b>
	H3.6Mr+3DHP	90.5*	80.9*	90.0*	85.6*	70.2*	93.0*	92.9*	83.8*	47.7*	85.0*

Table 3: Activitywise performance on MPI-INF-3DHP test set. We report the PCK, AUC and MPJPE. Higher PCK and AUC are better and lower MPJPE is better. (MPII) means pretrained on MPII dataset. “a” means adding background augmentation for training, “p” means Procrustes alignment, \* denote using groundtruth limb length.

We also test the model trained on Human3.6m that uses relative limb vector (to the length of the torso) instead of orientations. During the testing time, we rescaled each limb to its groundtruth length, only preserving the orientations. We show that it generalize worse than using orientation as representation.

### 4.3.2 Training on MPI-INF-3DHP Dataset

We finetune the network trained on the Human3.6m previous on this new dataset, without adding background augmentations. Results are shown in Table 3. Compared with [15, 14], which use extensive background and clothes augmentation, our model trained merely on the green-screen background can still outperform other methods. This shows that our method is more robust to background and can transfer well to different scenes.

## 4.4 Ablation Study

In this part, we evaluate different components of our approach on the Human3.6m. For simplicity, we use another protocol used in [9, 21, 60], where only one subject (S11) is used for testing. The estimated 3D pose is first aligned with a rigid transformation, which ruling out factors such as scale and rotation. We trained the 1-stack networks from scratch and 5-stack networks pretrained on the MPII dataset. Results are shown in table 4.



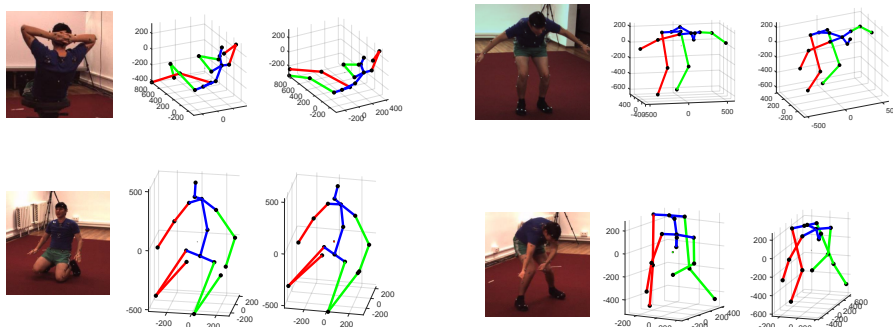


Figure 4: Qualitative results on Human3.6M dataset. In each example, the first one is the input image, the second one is our estimated pose and the third one is the groundtruth.

method	MPJPE	method	MPJPE
Chen[14]	82.72	Chen[14]+2D GT	57.50
1-stack	50.68	1-stack-fc	56.09
1-stack-len	86.32	1-stack-len-fc	81.51
5-stack-len	52.72	5-stack-len(rescaled)	41.95
5stack	<b>38.53</b>		

Table 4: The mean reconstruction errors on Human3.6M (S11). Errors are computed after Procrustes alignment. “fc” means sparse representation with a fully connected layer at the end. “len” means using the properly normalized limb vector instead of orientations. “Rescaled” means rescaled each bone to the groundtruth length during testing.

#### 4.4.1 Orientation

First, we demonstrate the advantage of estimating orientations compared with the original limb vector. We train the same network using the original limb vector properly normalized by the length of torso. As we can see, using the limb orientation representation can perform significantly better than considering bone length at the same time for both 1-stack and 5-stack networks (see 1-stack-len vs 1-stack and 5-stack-len vs 5-stack respectively). Even if we rescaled the bone vector to the groundtruth length preserving the orientation, the result is still worse than directly regressing the orientation. Furthermore, the 1-stack network using the orientation representation can even outperform 5-stack network with the original limb vector. This shows that decoupling the orientation and length of each limb is beneficial.

#### 4.4.2 Image to Prediction Association

In this part, we show the advantages of having the spatial associations between images and predictions. We train 1-stack networks from scratch, one with dense output as proposed, one with max pooling layer after the last several convolutional layers and the final convolutional layer replaced by a fully connected layer for direct regression. The dense representations achieves better results using orientations as representation. This show the benefits of our proposed image-to-prediction association. However, by considering limb length, the fully convolutional does not perform as well as the fully connected network. We conjecture that it

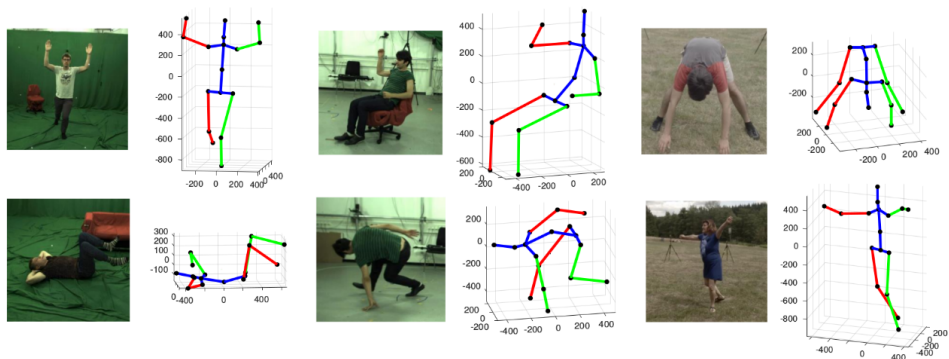


Figure 5: Qualitative results on the MPI-INF-3DHP test set. Our model trained only on green-screen background images (see column 1) can still perform well on novel scenes (see column 2 and 3).

may due to the relative small receptive field the network actually has. It also shows the necessary for using orientation for fully convolutional networks to achieve better performance.

#### 4.4.3 Robust to Bounding Box Jitter and Scale

Similar to [14], we carry out experiments on MPI-INF-3DHP dataset by jittering the bounding box at random in the range of  $\pm 40$  px. Since during the testing, we need to resize the bounding boxes to a fixed scale (256 in our experiments and 224 in [14]), we also add more noises by jittering the bounding box in the range of  $\pm 100$  px and rescale it by a factor of  $\pm 0.2$ . Our method is less sensitive to inaccurate bounding boxes.

	PCK	AUC
Mehta [14](40px)	70.1 ( $\downarrow \sim 6$ )	35.7
ours(40px)	80.8 $\pm$ 0.16 ( $\downarrow \sim 1$ )	44.3 $\pm$ 0.11
ours(100px+rescale)	77.9 ( $\downarrow \sim 4$ )	41.3 ( $\downarrow \sim 4$ )

## 5 Conclusion

In this paper, we propose a fully convolutional network that ties orientations with corresponding limb region to enhance the spatial relation between images and predictions. Our method is simple yet effective. Further experiments show that our model can generalize better to novel scenes and robust to inaccurate bounding boxes. In the future work, we expect adding additional constraints during both training and testing process would further improve the performance.

**Acknowledgement** This work is supported by the Intelligence Advanced Research Projects Activity (IARPA) via Department of Interior/ Interior Business Center (DOI/IBC) contract number D17PC00345. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon. Disclaimer: The views and conclusions contained herein are those of the authors and should

not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of IARPA, DOI/IBC, or the U.S. Government.

## References

- [1] Ijaz Akhter and Michael J Black. Pose-conditioned joint angle limits for 3d human pose reconstruction. In *CVPR*, 2015.
- [2] Mykhaylo Andriluka, Leonid Pishchulin, Peter Gehler, and Bernt Schiele. 2d human pose estimation: New benchmark and state of the art analysis. In *CVPR*, June 2014.
- [3] Zhe Cao, Tomas Simon, Shih-En Wei, and Yaser Sheikh. Realtime multi-person 2d pose estimation using part affinity fields. *CVPR*, 2017.
- [4] Ching-Hang Chen and Deva Ramanan. 3d human pose estimation= 2d pose estimation+ matching. *CVPR*, 2017.
- [5] Xiao Chu, Wanli Ouyang, Hongsheng Li, and Xiaogang Wang. Structured feature learning for pose estimation. In *CVPR*, 2016.
- [6] Xiao Chu, Wanli Ouyang, Xiaogang Wang, et al. Crf-cnn: Modeling structured information in human pose estimation. In *NIPS*, 2016.
- [7] Xiao Chu, Wei Yang, Wanli Ouyang, Cheng Ma, Alan Yuille, and Xiaogang Wang. Multi-context attention for human pose estimation. In *CVPR*, 2017.
- [8] Ronan Collobert, Koray Kavukcuoglu, and Clément Farabet. Torch7: A matlab-like environment for machine learning. In *BigLearn, NIPS Workshop*, 2011.
- [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*, 2018.
- [10] Ke Gong, Xiaodan Liang, Xiaohui Shen, and Liang Lin. Look into person: Self-supervised structure-sensitive learning and a new benchmark for human parsing. *CVPR*, 2017.
- [11] 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, jul 2014.
- [12] Sijin Li and Antoni B Chan. 3d human pose estimation from monocular images with deep convolutional neural network. In *ACCV*. Springer, 2014.
- [13] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *ECCV*, pages 740–755. Springer, 2014.
- [14] Julieta Martinez, Rayat Hossain, Javier Romero, and James J Little. A simple yet effective baseline for 3d human pose estimation. *ICCV*, 2017.
- [15] 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 *3D Vision (3DV), 2017 Fifth International Conference on*, .

- [16] 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. .
- [17] Francesc Moreno-Noguer. 3d human pose estimation from a single image via distance matrix regression. *CVPR*, 2017.
- [18] Alejandro Newell, Kaiyu Yang, and Jia Deng. Stacked hourglass networks for human pose estimation. In *ECCV*, 2016.
- [19] Georgios Pavlakos, Xiaowei Zhou, Konstantinos G Derpanis, and Kostas Daniilidis. Coarse-to-fine volumetric prediction for single-image 3d human pose. *CVPR*, 2017.
- [20] Varun Ramakrishna, Takeo Kanade, and Yaser Sheikh. Reconstructing 3d human pose from 2d image landmarks. In *ECCV*, 2012.
- [21] Grégory Rogez and Cordelia Schmid. Mocap-guided data augmentation for 3d pose estimation in the wild. In *NIPS*, 2016.
- [22] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision*, 115(3):211–252, Dec 2015.
- [23] Xiao Sun, Jiayang Shang, Shuang Liang, and Yichen Wei. Compositional human pose regression. *ICCV*, 2017.
- [24] Bugra Tekin, Pablo Márquez-Neila, Mathieu Salzmann, and Pascal Fua. Fusing 2d uncertainty and 3d cues for monocular body pose estimation. *ICCV*, 2017.
- [25] Denis Tome, Chris Russell, and Lourdes Agapito. Lifting from the deep: Convolutional 3d pose estimation from a single image. *CVPR*, 2017.
- [26] Alexander Toshev and Christian Szegedy. Deeppose: Human pose estimation via deep neural networks. In *CVPR*, 2014.
- [27] Chunyu Wang, Yizhou Wang, Zhouchen Lin, Alan L Yuille, and Wen Gao. Robust estimation of 3d human poses from a single image. In *CVPR*, 2014.
- [28] Shih-En Wei, Varun Ramakrishna, Takeo Kanade, and Yaser Sheikh. Convolutional pose machines. In *CVPR*, 2016.
- [29] Fangting Xia, Peng Wang, Liang-Chieh Chen, and Alan L Yuille. Zoom better to see clearer: Human and object parsing with hierarchical auto-zoom net. In *ECCV*, 2016.
- [30] Hashim Yasin, Umar Iqbal, Bjorn Kruger, Andreas Weber, and Juergen Gall. A dual-source approach for 3d pose estimation from a single image. In *CVPR*, 2016.
- [31] Xiaowei Zhou, Spyridon Leonardos, Xiaoyan Hu, and Kostas Daniilidis. 3d shape estimation from 2d landmarks: A convex relaxation approach. In *CVPR*, 2015.
- [32] Xiaowei Zhou, Menglong Zhu, Spyridon Leonardos, Konstantinos G Derpanis, and Kostas Daniilidis. Sparseness meets deepness: 3d human pose estimation from monocular video. In *CVPR*, 2016.

- [33] Xiaowei Zhou, Menglong Zhu, Georgios Pavlakos, Spyridon Leonardos, Kostantinos G Derpanis, and Kostas Daniilidis. Monocap: Monocular human motion capture using a cnn coupled with a geometric prior. *arXiv preprint arXiv:1701.02354*, 2017.
- [34] Xingyi Zhou, Xiao Sun, Wei Zhang, Shuang Liang, and Yichen Wei. Deep kinematic pose regression. In *ECCV 2016 Workshops*, 2016.
- [35] Xingyi Zhou, Qixing Huang, Xiao Sun, Xiangyang Xue, and Yichen Wei. Weakly-supervised transfer for 3d human pose estimation in the wild. *ICCV*, 2017.