

# Synthetic Occlusion Augmentation with Volumetric Heatmaps for the 2018 ECCV PoseTrack Challenge on 3D Human Pose Estimation

István Sárandi<sup>1</sup>, Timm Linder<sup>2</sup>, Kai O. Arras<sup>2</sup> and Bastian Leibe<sup>1</sup>

<sup>1</sup>Visual Computing Institute, RWTH Aachen University

{sarandi, leibe}@vision.rwth-aachen.de

<sup>2</sup>Robert Bosch GmbH, Corporate Research

{timm.linder, kaioliver.arras}@de.bosch.com

## Abstract

In this paper we present our winning entry at the 2018 ECCV PoseTrack Challenge on 3D human pose estimation. Using a fully-convolutional backbone architecture, we obtain volumetric heatmaps per body joint, which we convert to coordinates using soft-argmax. Absolute person center depth is estimated by a 1D heatmap prediction head. The coordinates are back-projected to 3D camera space, where we minimize the L1 loss. Key to our good results is the training data augmentation with randomly placed occluders from the Pascal VOC dataset. In addition to reaching first place in the Challenge, our method also surpasses the state-of-the-art on the full Human3.6M benchmark when considering methods that use no extra pose datasets in training. Code for applying synthetic occlusions is available at <https://github.com/isarandi/synthetic-occlusion>.

## 1. Introduction

The 3D part of the 2018 ECCV PoseTrack Challenge invited participants to tackle the following task. Given an uncropped, static RGB image containing a single person, estimate the position of  $J$  body joints in 3D camera space, relative to the root (pelvis) joint position. Predicting human poses in 3D space has several important applications, such as human-robot collaboration and virtual reality.

## 2. Related Work

**3D Human Pose Estimation.** State-of-the-art 3D pose estimation methods are based on deep convolutional neural networks. We recommend Sarafianos *et al.*'s survey [22] for an overview of methods. Recently, building on experience gained from 2D human pose estimation (*e.g.*, [16]), heatmap-like methods have been introduced for 3D pose es-




Figure 1. Examples of synthetic occlusions with Pascal VOC objects (geometric and color augmentations not depicted).

timation with promising results. This includes volumetric heatmaps [20][27][13], marginal heatmaps [17] and location maps [15].

**Occlusion Augmentation.** Erasing or pasting over parts of an image has been successfully used as data augmentation in image classification, object detection, person re-identification [30][2][4][3][6], and facial landmark localization [29]. Ke *et al.* [11] augment images for 2D pose estimation by copying background patches over some of the body joints. We have recently found that such techniques are also very effective for 3D pose estimation [23].

## 3. Dataset

The dataset in this challenge is a subset of Human3.6M [10][9], with 35,832 training, 19,312 validation and 24,416 test examples. There are a few important differences compared to the full benchmark. First, the Challenge version lacks person bounding boxes and camera intrinsics as input. Second, the ground truth labels are more restricted, consisting only of camera-space 3D joint coordinates after




Figure 2. Overview of our architecture.

subtraction of the root joint. Image-space joint coordinates are not available either.

## 4. Approach

We present a modified version of the method we recently used for studying occlusion-robustness in 3D pose estimation [23], extending it to handle the above-mentioned differences.

**Image Preprocessing.** We obtain person bounding boxes using the YOLOv3 detector [21]. Treating the original camera’s focal length  $f$  as a global hyperparameter, we reproject the image to be centered on the person box, at a scale where the larger side of the box fills 90% of the output.

**Backbone Network.** We feed the cropped and zoomed image ( $256 \times 256$  px) into a fully-convolutional backbone network (ResNet v2-50 [7][24]). We directly obtain volumetric heatmaps from the backbone net by adding a  $1 \times 1$  convolutional layer on the last spatial feature map of the backbone, producing  $J \cdot D$  output channels. The resulting tensor is reshaped to yield  $J$  volumes, one per body joint, each with depth  $D$ .

**Volumetric Heatmaps.** We follow Pavlakos *et al.* in the interpretation of the volumetric heatmap’s axes [20]: X and Y correspond to image space and the depth axis to camera space, relative to the person center. Relative depths are not sufficient, however, when back-projecting from image to camera space. Pavlakos *et al.* optimize the root joint depth in post-processing, based on bone-length priors. By contrast, we predict it using a second prediction head on the backbone net (see Fig. 2). This outputs a 1D heatmap discretized to 32 units, representing a 10 meter range in front of the camera.

**Soft-Argmax.** We extract coordinate predictions from the heatmaps using soft-argmax [12][18]. Since this operation is differentiable, there is no need to provide ground-truth heatmaps at training time [27]. Instead, the loss can be computed deeper in the network and backpropagated through the soft-argmax operation. Soft-argmax also reduces the quantization errors inherent in hard argmax and gives fine-grained, continuous results without requiring memory-expensive, high-resolution heatmaps [27]. Indeed,

we use a heatmap resolution as low as  $16^3$  for the results presented in this paper.

**Camera Intrinsic.** Having predicted image coordinates  $x_i, y_i$ , depth coordinates  $\Delta Z_i$  relative to the person center and the absolute depth  $Z^*$  of the person center by soft-argmax, we now need camera intrinsics to move from image space to 3D camera space. As mentioned earlier, the original camera’s focal length  $f$  is treated as a hyperparameter, and we must also take into account the zooming factor  $s$  applied in preprocessing.

To avoid the need for precise hyperparameter tuning of  $f$ , we learn an additional, input-independent corrective factor  $c$  for the focal length during training, to achieve better alignment of image and heatmap locations. Denoting the image height and width as  $H$  and  $W$ , back-projection is performed as

$$\begin{pmatrix} X_i \\ Y_i \\ Z_i \end{pmatrix} = (Z^* + \Delta Z_i) \begin{pmatrix} fsc & 0 & W/2 \\ 0 & fsc & H/2 \\ 0 & 0 & 1 \end{pmatrix}^{-1} \begin{pmatrix} x_i \\ y_i \\ 1 \end{pmatrix}. \quad (1)$$

**Loss.** After subtracting the root joint coordinates, we compute the  $L^1$  loss in the original camera space w.r.t. the provided root-relative ground truth. No explicit heatmap loss is used. Since all above operations are differentiable the whole network can be trained end-to-end.

**Data Augmentation.** In our recent study on the occlusion-robustness of 3D pose estimation [23], we found that augmenting training images with synthetic occlusions acts as an effective regularizer. Starting with the objects in the Pascal VOC dataset [5], we filter out persons, segments labeled as *difficult* or *truncated* and segments with area below 500 px, leaving 2638 objects. With probability  $p_{occ}$ , we paste a random number (between 1 and 8) of these objects at random locations in each frame. We also apply standard geometric augmentations (scaling, rotation, translation, horizontal flip) and appearance distortions (blurs and color manipulations). At test time only horizontal flipping augmentation is used.

**Training Details.** The backbone net is initialized with ImageNet-pretrained weights from [24]. We train the final method for 410 epochs on the union of the training and val-

	Direct	Discuss	Eat	Greet	Phone	Pose	Purch.	Sit	SitD	Smoke	Photo	Wait	Walk	WalkD	WalkT	Avg
Zhu	58	59	64	62	65	60	68	77	92	65	68	62	60	70	59	66
Rhodin	51	53	58	52	64	53	67	94	132	65	64	57	53	67	53	66
Zhou	52	56	55	51	57	53	64	73	81	61	60	57	49	61	53	59
Park	53	52	52	53	55	55	54	71	84	56	60	58	51	64	57	58
Shen	53	54	54	52	56	55	58	70	78	60	59	57	48	61	56	58
Pavlakos	44	46	50	47	56	47	52	63	70	54	54	48	46	58	46	52
Sun	<b>38</b>	43	46	41	46	<b>40</b>	49	65	73	48	49	43	38	52	<b>38</b>	47
<b>Ours</b>	<b>38</b>	<b>40</b>	<b>43</b>	<b>40</b>	<b>43</b>	<b>40</b>	<b>47</b>	<b>58</b>	<b>64</b>	<b>43</b>	<b>48</b>	<b>42</b>	<b>36</b>	<b>50</b>	<b>38</b>	<b>45</b>

Table 1. Mean per joint position errors achieved by participants of the 2018 ECCV PoseTrack Challenge on 3D human pose estimation [1], on a subset of the Human3.6M dataset. In contrast to our method, some participants used extra 2D pose datasets in training (in accordance with the challenge rules).

	Direct	Discuss	Eat	Greet	Phone	Pose	Purch.	Sit	SitD	Smoke	Photo	Wait	Walk	WalkD	WalkT	Avg
* Zhou (2017) [31]	54.8	60.7	58.2	71.4	62.0	65.5	53.8	<b>55.6</b>	75.2	111.6	64.2	66.0	51.4	63.2	55.3	64.9
* Martinez (2017) [14]	51.8	56.2	58.1	59.0	69.5	55.2	58.1	74.0	94.6	62.3	78.4	59.1	65.1	49.5	52.4	62.9
* Sun (2017) [26]	52.8	54.8	54.2	54.3	61.8	53.1	53.6	71.7	86.7	61.5	67.2	53.4	47.1	61.6	53.4	59.1
* Pavlakos (2018) [19]	48.5	54.4	54.4	52.0	59.4	49.9	52.9	65.8	71.1	56.6	65.3	52.9	60.9	<b>44.7</b>	47.8	56.2
* Luvizon (single-crop, 2018) [13]	51.5	53.4	49.0	52.5	53.9	50.3	54.4	63.6	73.5	55.3	61.9	50.1	46.0	60.2	51.0	55.1
* Luvizon (multi-crop, 2018) [13]	49.2	51.6	<b>47.6</b>	50.5	51.8	48.5	51.7	61.5	70.9	53.7	60.3	48.9	44.4	57.9	48.9	53.2
* Sun (2018) [27]	<b>47.5</b>	<b>47.7</b>	49.5	<b>50.2</b>	<b>51.4</b>	<b>43.8</b>	<b>46.4</b>	58.9	<b>65.7</b>	<b>49.4</b>	<b>55.8</b>	<b>47.8</b>	<b>38.9</b>	49.0	43.8	<b>49.6</b>
Tekin (2016) [28]	102.4	147.7	88.8	125.4	118.0	112.4	129.2	138.9	224.9	118.4	182.7	138.8	55.1	126.3	65.8	125.0
Zhou (2016) [33]	87.4	109.3	87.1	103.2	116.2	106.9	99.8	124.5	199.2	107.4	139.5	118.1	79.4	114.2	97.7	113.0
Xingyi (2016) [32]	91.8	102.4	97.0	98.8	113.4	90.0	93.8	132.2	159.0	106.9	125.2	94.4	79.0	126.0	99.0	107.3
Sun (2017) [26]	90.2	95.5	82.3	85.0	87.1	87.9	93.4	100.3	135.4	91.4	94.5	87.3	78.0	90.4	86.5	92.4
Pavlakos (2017) [20]	67.4	72.0	66.7	69.1	72.0	65.0	68.3	83.7	96.5	71.7	77.0	65.8	59.1	74.9	63.2	71.9
Sun (2018) [27]	63.8	64.0	56.9	64.8	62.1	59.8	60.1	71.6	91.7	60.9	70.4	65.1	<i>51.3</i>	63.2	55.4	64.1
<b>Ours (no occlusion augmentation)</b>	63.3	65.5	56.0	62.1	64.0	60.7	64.8	76.7	93.0	63.3	69.7	62.0	68.8	61.3	54.1	65.7
<b>Ours (full)</b>	<i>49.1</i>	<i>54.6</i>	<i>50.4</i>	<i>50.7</i>	<i>54.8</i>	<i>47.4</i>	<i>50.1</i>	<i>67.5</i>	<i>78.4</i>	<i>53.1</i>	<i>57.4</i>	<i>50.7</i>	<i>54.0</i>	<i>46.1</i>	<b>40.1</b>	<i>54.2</i>

Table 2. Mean per joint position error on the full Human3.6M dataset. Results marked with an asterisk (\*) were achieved using extra 2D pose dataset(s) in training. Boldface indicates the overall best results, while italic indicates the best when using no extra 2D pose datasets.

ication set using the Adam optimizer and cyclical (triangular) learning rates [25]. Our final challenge predictions were produced using a snapshot ensemble [8], averaging the predictions of snapshots taken at the last three learning-rate-minima of the cyclical schedule. We used  $f = 1500$  and  $p_{occ} = 0.5$  for the submission.

## 5. Results

The evaluation metric is the mean per joint position error (MPJPE) over all joints after subtraction of the root joint position. Our method achieves best results for all actions, even ahead of methods using extra 2D pose datasets in training (see Table 1). The margin is largest for the actions *Sitting* and *Sitting Down*, showing that our method is more robust to the presence of a chair, which is the only occluding object in the Human3.6M dataset.

**Effect of Occlusion Augmentation.** Fig. 3 shows how synthetic occlusion augmentation improves results on the Challenge validation set as we vary the probability  $p_{occ}$  of applying occlusion augmentation to each frame. Augmenting just 10% of the images already improves MPJPE by 8.2 mm and improvements continue to about  $p_{occ} = 70\%$ , after which performance is only influenced slightly.




Figure 3. Effect of the per-frame probability of occlusion augmentation ( $p_{occ}$ ) evaluated on the Challenge validation set.

**Full Human3.6M Benchmark.** For comparison with prior work, we train and evaluate our method on the full Human3.6M benchmark as well (Table 2). Here we use the bounding boxes and camera intrinsics provided with the dataset and minimize the  $L^1$  loss computed on the absolute (*i.e.* non-root-relative) coordinates in camera space for 40 epochs. The person center depth  $Z^*$  is estimated as described in Section 4. We follow the common protocol of training on five subjects (S1, S5, S6 S7, S8) and evaluating on two (S9, S11), without Procrustes alignment. We use no snapshot ensembling here, for better comparability.

The occlusion probability  $p_{occ}$  is set to 1. Our method outperforms all prior work on Human3.6M when no additional pose datasets are used for training.

## 6. Conclusion

We have presented an architecture and data augmentation method for 3D human pose estimation and have shown that it outperforms other methods both by achieving first place in the 2018 ECCV PoseTrack Challenge and by surpassing the state-of-the-art on the full benchmark among methods using no additional pose datasets in training.

**Acknowledgments.** This project has been funded by a grant from the Bosch Research Foundation, by ILIAD (H2020-ICT-2016-732737) and by ERC Consolidator Grant DeeViSe (ERC-2017-COG-773161).

## References


- [1] 3D pose estimation rankings of the 2018 ECCV PoseTrack Challenge. <http://archive.today/2018.09.12-150455/http://vision.imar.ro/human3.6m/ranking.php>. Accessed: 2018-09-12.
- [2] T. DeVries and G. W. Taylor. Improved regularization of convolutional neural networks with cutout. *arXiv:1708.04552*, 2017.
- [3] N. Dvornik, J. Mairal, and C. Schmid. Modeling visual context is key to augmenting object detection datasets. In *ECCV*, 2018.
- [4] D. Dwibedi, I. Misra, and M. Hebert. Cut, paste and learn: Surprisingly easy synthesis for instance detection. In *ICCV*, 2017.
- [5] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The PASCAL Visual Object Classes Challenge 2012 (VOC2012) Results. <http://www.pascal-network.org/challenges/VOC/voc2012/workshop/index.html>.
- [6] G. Georgakis, A. Mousavian, A. C. Berg, and J. Kosecka. Synthesizing training data for object detection in indoor scenes. *arXiv:1702.07836*, 2017.
- [7] K. He, X. Zhang, S. Ren, and J. Sun. Identity mappings in deep residual networks. In *ECCV*, 2016.
- [8] G. Huang, Y. Li, G. Pleiss, Z. Liu, J. E. Hopcroft, and K. Q. Weinberger. Snapshot ensembles: Train 1, get M for free. *arXiv:1704.00109*, 2017.
- [9] C. Ionescu, F. Li, and C. Sminchisescu. Latent structured models for human pose estimation. In *ICCV*, 2011.
- [10] C. Ionescu, D. Papava, V. Olaru, and C. Sminchisescu. Human3.6M: Large scale datasets and predictive methods for 3D human sensing in natural environments. *PAMI*, 2014.
- [11] L. Ke, M.-C. Chang, H. Qi, and S. Lyu. Multi-scale structure-aware network for human pose estimation. *arXiv:1803.09894*, 2018.
- [12] S. Levine, C. Finn, T. Darrell, and P. Abbeel. End-to-end training of deep visuomotor policies. *JMLR*, 17(1):1334–1373, 2016.
- [13] D. C. Luvizon, D. Picard, and H. Tabia. 2D/3D pose estimation and action recognition using multitask deep learning. In *CVPR*, 2018.
- [14] J. Martinez, R. Hossain, J. Romero, and J. J. Little. A simple yet effective baseline for 3D human pose estimation. In *ICCV*, 2017.
- [15] D. Mehta et al. Vnect: Real-time 3d human pose estimation with a single RGB camera. *ACM Trans. Graphics*, 36(4):44, 2017.
- [16] A. Newell, K. Yang, and J. Deng. Stacked hourglass networks for human pose estimation. In *ECCV*, 2016.
- [17] A. Nibali, Z. He, S. Morgan, and L. Prendergast. 3D human pose estimation with 2D marginal heatmaps. *arXiv:1806.01484*, 2018.
- [18] A. Nibali, Z. He, S. Morgan, and L. Prendergast. Numerical coordinate regression with convolutional neural networks. *arXiv:1801.07372*, 2018.
- [19] G. Pavlakos, X. Zhou, and K. Daniilidis. Ordinal depth supervision for 3D human pose estimation. In *CVPR*, 2018.
- [20] G. Pavlakos, X. Zhou, K. G. Derpanis, and K. Daniilidis. Coarse-to-fine volumetric prediction for single-image 3D human pose. In *CVPR*, 2017.
- [21] J. Redmon and A. Farhadi. YOLOv3: An incremental improvement. *arXiv:1804.02767*, 2018.
- [22] N. Sarafianos, B. Boteanu, B. Ionescu, and I. A. Kakadiaris. 3D human pose estimation: A review of the literature and analysis of covariates. *CVIU*, 152:1–20, 2016.
- [23] I. Sárándi, T. Linder, K. O. Arras, and B. Leibe. How robust is 3D human pose estimation to occlusion? In *IROS Workshops*, 2018. *arXiv:1808.09316*.
- [24] N. Silberman and S. Guadarrama. TensorFlow-Slim image classification model library. <https://github.com/tensorflow/models/tree/master/research/slim>, 2016.
- [25] L. N. Smith. Cyclical learning rates for training neural networks. In *WACV*, 2017.
- [26] X. Sun, J. Shang, S. Liang, and Y. Wei. Compositional human pose regression. In *ICCV*, 2017.
- [27] X. Sun, B. Xiao, F. Wei, S. Liang, and Y. Wei. Integral human pose regression. In *ECCV*, 2018.
- [28] B. Tekin, A. Rozantsev, V. Lepetit, and P. Fua. Direct prediction of 3D body poses from motion compensated sequences. In *CVPR*, 2016.
- [29] K. Yuen and M. M. Trivedi. An occluded stacked hourglass approach to facial landmark localization and occlusion estimation. *IEEE Trans. Intel. Veh.*, 2(4):321–331, 2017.
- [30] Z. Zhong, L. Zheng, G. Kang, S. Li, and Y. Yang. Random erasing data augmentation. *arXiv:1708.04896*, 2017.
- [31] X. Zhou, Q. Huang, X. Sun, X. Xue, and Y. Wei. Towards 3D human pose estimation in the wild: a weakly-supervised approach. In *ICCV*, 2017.
- [32] X. Zhou, X. Sun, W. Zhang, S. Liang, and Y. Wei. Deep kinematic pose regression. In *ECCV*, 2016.
- [33] X. Zhou, M. Zhu, S. Leonardos, K. G. Derpanis, and K. Daniilidis. Sparseness meets deepness: 3D human pose estimation from monocular video. In *CVPR*, 2016.

# Synthetic Occlusion Augmentation for 3D Human Pose Estimation with Volumetric Heatmaps


 István Sáránci<sup>1</sup>, Timm Linder<sup>2</sup>, Kai O. Arras<sup>2</sup>, Bastian Leibe<sup>1</sup>
<sup>1</sup>Visual Computing Institute, RWTH Aachen University – Aachen, Germany

<sup>2</sup>Robert Bosch GmbH, Corporate Research – Stuttgart, Germany

## PoseTrack Challenge 2018 – 3D Task



**Input image**  
 Uncropped, static RGB image of a single person



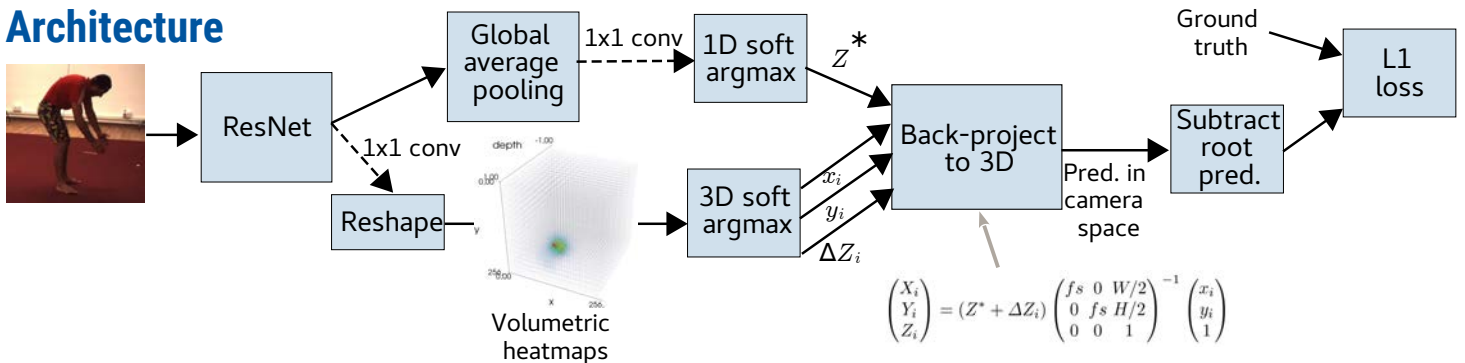
**Output 3D skeleton**  
 17 body joints in 3D camera space relative to the root (pelvis) joint

```
[[ 0.0 0.0 0.0]
 [ -96.9 -21.1 163.4]
 [ 43.0 456.4 150.3]
 [ 42.6 902.6 249.0]
 [ 96.0 22.1 -102.8]
 [ 91.0 508.3 -100.2]
 [ 118.4 953.7 13.0]
 [ 3.1 -262.6 13.2]
 [ -36.6 -502.0 -72.2]
 [ -96.7 -541.9 -162.8]
 [ -88.0 -651.5 -140.0]
 [ 85.4 -439.5 -131.1]
 [ 278.0 -206.0 -121.3]
 [ 367.3 28.3 -184.0]
 [ -139.0 465.2 28.8]
 [ -406.6 322.4 42.9]
 [ -414.1 -255.1 -202.1]]
```

## Our Approach

- Detect person with YOLOv3, then zoom & crop
  - Predict **volumetric body joint heatmaps** directly, with a fully-convolutional backbone (ResNet-50v2)
  - Predict person center depth with a **1D heatmap head**
  - Obtain 3D points with **soft-argmax** and camera **back-projection**
  - Minimize the **L1 loss** after subtracting root joint
- 
- Achieved **first place** in the Challenge
  - **No additional pose datasets** used for training
  - **High frame rate** inference (204 fps, excl. detection) on Titan X GPU

## Architecture



## Occlusion Augmentation at Training



**2638 occluder objects from Pascal VOC**  
 Filter out 'person', 'truncated', 'difficult' and small object segments

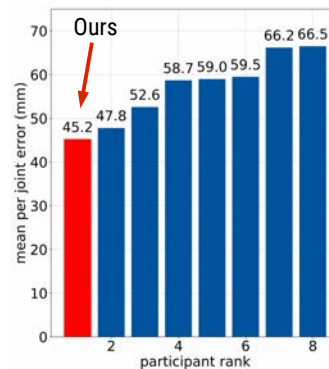


**Augmented inputs with pasted occluders**  
 Applied with 50% probability, 1–8 objects, at random scale, at random position



**Why not some simpler geometric shapes?**  
 We found them less effective in our recent occlusion-robustness study<sup>[8]</sup>

## Quantitative Results



**Final challenge results**

Method	Extra pose data in training?	
	no	yes
Sun (ICCV'17) [3]	92.4	59.1
Martinez (ICCV'17) [4]	–	62.9
Zhou (ICCV'17) [5]	–	55.9
Pavlakos (CVPR'18) [6]	71.9	56.2
Sun (ArXiv) [7]	64.1	<b>49.6</b>
Ours (no occlusion augm.)	65.7	–
Ours (full)	<b>55.4</b>	–

**Comparison on the full Human3.6M<sup>[1][2]</sup> benchmark**  
 MPJPE, trained on subjects S1, S5, S6, S7, S8; tested on S9, S11

## References

- [1] Ionescu, C., Li, F., Sminchisescu, C.: Latent structured models for human pose estimation. ICCV (2011)
- [2] Ionescu, C., Papava, D., Olaru, V., Sminchisescu, C.: Human3.6M: Large scale datasets and predictive methods for 3D human sensing in natural environments. PAMI (2014)
- [3] Sun, X., Shang, J., Liang, S., Wei, Y.: Compositional human pose regression. ICCV (2017).
- [4] Martinez, J., Hossain, R., Romero, J., Little, J.J.: A simple yet effective baseline for 3D human pose estimation. ICCV (2017)
- [5] Zhou, X., Huang, Q., Sun, X., Xue, X., Wei, Y.: Towards 3D human pose estimation in the wild: a weakly-supervised approach. ICCV (2017)
- [6] Pavlakos, G., Zhou, X., Daniilidis, K.: Ordinal depth supervision for 3D human pose estimation. CVPR (2018)
- [7] Sun, X., Xiao, B., Liang, S., Wei, Y.: Integral human pose regression. ArXiv:1711.08229 (2017)
- [8] Sáránci, I., Linder, T., Arras, K.O., Leibe, B.: How robust is 3D human pose estimation to occlusion? ArXiv:1808.09316 (2018)

## Qualitative Results



## Test set (final winning model, unknown ground truth)

