

Compressed Volumetric Heatmaps for Multi-Person 3D Pose Estimation

Matteo Fabbri^{1*} Fabio Lanzi¹ Simone Calderara¹ Stefano Alletto² Rita Cucchiara¹

¹University of Modena and Reggio Emilia

{name.surname}@unimore.it

²Panasonic R&D Company of America

{name.surname}@us.panasonic.com

Abstract

In this paper we present a novel approach for bottom-up multi-person 3D human pose estimation from monocular RGB images. We propose to use high resolution volumetric heatmaps to model joint locations, devising a simple and effective compression method to drastically reduce the size of this representation. At the core of the proposed method lies our Volumetric Heatmap Autoencoder, a fully-convolutional network tasked with the compression of ground-truth heatmaps into a dense intermediate representation. A second model, the Code Predictor, is then trained to predict these codes, which can be decompressed at test time to re-obtain the original representation. Our experimental evaluation shows that our method performs favorably when compared to state of the art on both multi-person and single-person 3D human pose estimation datasets and, thanks to our novel compression strategy, can process full-HD images at the constant runtime of 8 fps regardless of the number of subjects in the scene. Code and models are publicly available.

1. Introduction

Human Pose Estimation (HPE) has seen significant progress in recent years, mainly thanks to deep Convolutional Neural Networks (CNNs). Best performing methods on 2D HPE are all leveraging heatmaps to predict body joint locations [3, 49, 43]. Heatmaps have also been extended for 3D HPE, showing promising results in single person contexts [38, 29, 41].

Despite their good performance, these methods do not easily generalize to multi-person 3D HPE, mainly because of their high demands for memory and computation. This drawback also limits the resolution of those maps, that have to be kept small, leading to quantization errors. Using larger volumetric heatmaps can address those issues, but at the cost of extra storage, computation and training complexity.




Figure 1: Examples of 3D poses estimated by our LoCO approach. Close-ups show that 3D poses are correctly computed even in very complex and articulated scenarios

In this paper, we propose a simple solution to the aforementioned problems that allows us to directly predict high-resolution volumetric heatmaps while keeping storage and computation small. This new solution enables our method to tackle multi-person 3D HPE using heatmaps in a single-shot bottom-up fashion. Moreover, thanks to our high-resolution output, we are able to produce fine-grained absolute 3D predictions even in single person contexts. This allows our method to achieve state of the art performance on the most popular single person benchmark [11].

The core of our proposal relies on the creation of an alternative ground-truth representation that preserves the most informative content of the original ground-truth but reduces its memory footprint. Indeed, this new compressed representation is used as the target ground-truth during our network training. We named this solution LoCO, *Learning on Compressed Output*.

By leveraging on the analogy between compression and dimensionality reduction on sparse signals [47, 39, 1], we empirically follow the intuition that 3D body poses can be represented in an alternative space where data redundancy

* Work done while interning at Panasonic R&D Company of America

is exploited towards a compact representation. This is done by minimizing the loss of information while keeping the spatial nature of the representation, a task for which convolutional architectures are particularly suitable. Concurrently w.r.t. our proposal, compression-based approaches have been effectively used for both dataset distillation and input compression [48, 46] but, to the best of our knowledge, this is the first time they are applied to ground truth remapping. For this purpose, deep self-supervised networks such as autoencoders represent a natural choice for searching, in a data-driven way, for an intermediate representation.

Specifically, our HPE pipeline consists of two modules: at first, the pretrained *Volumetric Heatmap Autoencoder* is used to obtain a smaller/denser representation of the volumetric heatmaps. These “codes” are then used to supervise the *Code Predictor*, which aims at estimating multiple 3D joint locations from a monocular RGB input.

To summarize, the novel aspects of our proposal are:

- We propose a simple and effective method that maps high-resolution volumetric heatmaps to a compact and more tractable representation. This saves memory and computational resources while keeping most of the informative content.
- This new data representation enables the adoption of volumetric heatmaps to tackle multi-person 3D HPE in a bottom-up fashion, an otherwise intractable problem. Experiments on both real [12] and simulated environments [8] (see Fig. 1) show promising results even in 100 meters wide scenes with more than 50 people. Our method only requires a single forward pass and can be applied with constant running time regardless of the number of subjects in the scene.
- We further demonstrate the generalization capabilities of LoCO by applying it to a single person context. Our fine-grained predictions establish a new state of the art on Human3.6m [11] among bottom-up methods.

2. Related Work

Single-Person 3D HPE Single person 3D HPE from a monocular camera has become extremely popular in the last few years. Literature can be classified into three different categories: (i) approaches that first estimate 2D joints and then project them to 3D space, (ii) works that jointly estimate 2D and 3D poses, (iii) methods that learn the 3D pose directly from the RGB image.

The majority of works on single person 3D HPE first compute 2D poses and leverages them to estimate 3D poses, either using off-the-shelf 2D HPE methods [15, 10, 19, 20,

2, 24, 4] or by having a dedicated module in the 3D HPE pipeline [26, 28, 16, 51].

Joint learning of 2D and 3D pose is also shown to be beneficial [22, 6, 50, 54, 44, 27, 14, 30], often in conjunction with large-scale datasets that only provide 2D pose ground-truth and exploiting anatomical or structure priors.

Finally, recent works estimate 3D pose information directly [38, 29, 41, 18, 25, 34, 35]. Among these, Pavlakos *et al.* [29] were the first to propose a fine discretization of the 3D space around the target by learning a coarse-to-fine prediction scheme in an end to end fashion.

Multi-Person 3D HPE To the best of our knowledge, very few works tackle multi-person 3D HPE from monocular images. We can categorize them into two classes: top-down and bottom-up approaches.

Top-down methods first identify bounding boxes likely to contain a person using third party detectors and then perform single-person HPE for each person detected. Among them, Rogez *et al.* [37] classifies bounding boxes into a set of K-poses. These poses are scored by a classifier and refined using a regressor. The method implicitly reasons using bounding boxes and produces multiple proposals per subject that need to be accumulated and fused. Zanfir *et al.* [52] combine a single person model that incorporates feed-forward initialization and semantic feedback, with additional constraints such as ground plane estimation, mutual volume exclusion, and joint inference. Dabral *et al.* [6], instead, propose a two-staged approach that first estimates the 2D keypoints in every Region of Interest and then lifts the estimated keypoints to 3D. Finally, Moon *et al.* [23] predict absolute 3D human root localization, and root-relative 3D single-person for each person independently. However, these methods heavily rely on the accuracy of the people detector and do not scale well when facing scenes with dozens of people.

In contrast to top-down approaches, bottom-up methods produce multi-person joint locations in a single shot, from which the 3D pose can be inferred even under strong occlusions. Mehta *et al.* [21], predict 2D and 3D poses for all subjects in a single forward pass regardless of the number of people in the scene. They exploit occlusion-robust pose-maps that store 3D coordinates at each joint 2D pixel location. However, their 3D pose read-out strategy strongly depends on the 2D pose output which makes it limited by the accuracy of the 2D module. Their method also struggles to resolve scenes with multiple overlapping people, due to the missing 3D reasoning in their joint-to-person association process. Zanfir *et al.* [53], on the other hand, utilize a multi-task deep neural network where the person grouping problem is formulated as an integer program based on learned body part scores parameterized by both 2D and 3D information. Similarly to the latter, our method directly learns a

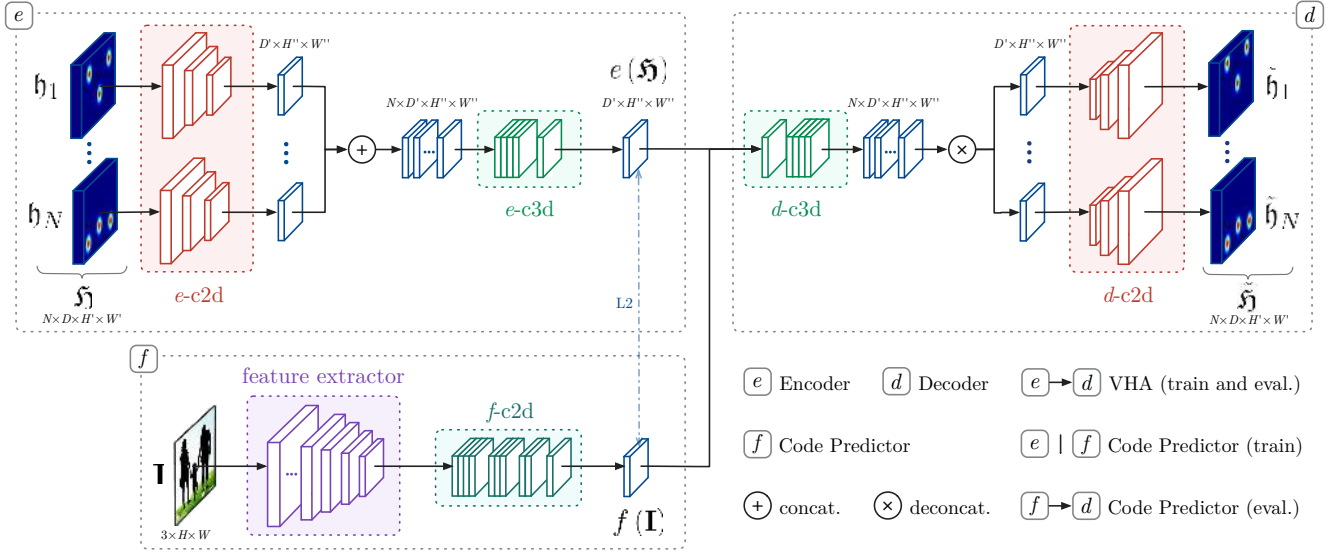


Figure 2: Schematization of the proposed LoCO pipeline. At training time, the Encoder e produces the compressed volumetric heatmaps $e(\mathfrak{h})$ which are used as ground truth from the Code Predictor f . At test time, the intermediate representation $f(I)$ computed by the Code Predictor is fed to the Decoder d for the final output. In our case, $H' = H/8$ and $W' = W/8$

mapping from image features to 3D joint locations, with no need of explicit bounding box detections or 2D proxy poses, while simultaneously being robust to heavy occlusions and multiple overlapping people.

Multi-Person 3D Pose Representation In a top-down framework, the simplest 3D pose representation can be expressed by a vector of joints. By casting 3D HPE as a coordinate regression task, Rogez *et al.* [37] and Zanfir *et al.* [52] indeed utilize x , y , z coordinates of the human joints w.r.t. a known root location. On the other hand, bottom-up approaches require a representation whose coding does not depend on the number of people (e.g. an image map). Among the most recent methods, Mehta *et al.* [21] and Zanfir *et al.* [53] both utilize a pose representation composed by joint-specific feature channels storing the 3D coordinate x , y , or z at the joint/limb 2D pixel location. This representation, however, suffers when multiple overlapping people are present in the scene. In contrast to all these approaches, we adopted the volumetric heatmap representation proposed by Pavlakos *et al.* [29], overcoming all the limitations that arise when facing a multi-person context.

3. Proposed Method

The following subsections summarize the key elements of LoCO. Section 3.1 gives a preliminary definition of the chosen volumetric heatmap representation and elaborates on its merits. Section 3.2 illustrates our proposed data mapping which addresses the high dimensional nature of the volumetric heatmaps by producing a compact and more tractable

representation. Next, in Section 3.3, we describe how our strategy can be easily exploited to effectively tackle the problem of multi-person 3D HPE in a single-shot bottom-up fashion. Finally, Section 3.4 illustrates our simple refining approach that prevents poses from being implausible.

3.1. Volumetric Heatmaps

By considering a voxelization of the RGB-D volumetric space [7, 29], we refer as a volumetric heatmap, \mathfrak{h} , the 3D confidence map with size $D \times H \times W$, where D represents the depth dimension (appropriately quantized), while H and W represent the height and width of the image plane respectively. Given the body joint j with pseudo-3D coordinates $\mathbf{u}_j = (u_{1,j}, u_{2,j}, u_{3,j})$, where $u_{1,j} \in \{1, \dots, D\}$ is the quantized distance of joint j from the camera, and $u_{2,j} \in \{1, \dots, H\}$ and $u_{3,j} \in \{1, \dots, W\}$ are respectively the row and column indexes of its pixel on the image plane, the value of \mathfrak{h}_j at a generic location \mathbf{u} is obtained by centering a fixed variance Gaussian in u_j :

$$\mathfrak{h}_j(\mathbf{u}) = e^{-\frac{\|\mathbf{u}-\mathbf{u}_j\|^2}{\sigma^2}} \quad (1)$$

In a multi-person context, in the same image we can simultaneously have several joints of the same kind (e.g. “left ankle”), one for each of the K different people in the image. In this case we aggregate those K volumetric heatmaps $\mathfrak{h}_j^{(k)}$, into a single heatmap \mathfrak{h}_j with a max operation:

$$\mathfrak{h}_j(\mathbf{u}) = \max_k \{\mathfrak{h}_j^{(k)}(\mathbf{u})\} \quad (2)$$

Finally, considering N different types of joint and K

block	layer	in ch.	out ch.	stride
<i>e-c2d</i>	Conv2D + ReLU	D	D/d_1	s_1
	Conv2D + ReLU	D/d_1	D/d_2	s_2
	Conv2D + ReLU	D/d_2	D/d_3	s_2
<i>e-c3d</i>	Conv3D + ReLU	N	4	1
	Conv3D + ReLU	4	1	1

Table 1: Structure of the encoder part of the Volumetric Heatmap Autoencoder (VHA). The decoder is not shown as it is perfectly mirrored to the encoder. *VHAv1*: $(d_1, d_2, d_3) = (1, 2, 2)$ and $(s_1, s_2, s_3) = (1, 2, 1)$; for *VHAv2*: $(d_1, d_2, d_3) = (2, 4, 4)$ and $(s_1, s_2, s_3) = (2, 2, 1)$; *VHAv3*: $(d_1, d_2, d_3) = (2, 4, 8)$ and $(s_1, s_2, s_3) = (2, 2, 2)$

people, we have a set of N volumetric heatmaps (each associated with a joint type), $\mathfrak{h} = \{h_j, j = 1, \dots, N\}$, resulting from the aggregation of the individual heatmaps of the K people in the scene. Note that, given pseudo-3D coordinates $\mathbf{u} = (u_1, u_2, u_3)$ and the camera intrinsic parameters, i.e. focal length $f = (f_x, f_y)$ and principal point (c_x, c_y) , the corresponding 3D coordinates $\mathbf{x} = (x, y, z)$ in the camera reference system can be retrieved by directly applying the equations of the pinhole camera model.

The benefit of choosing a volumetric heatmap representation over a direct 3D coordinate regression is that it casts the highly non-linear problem to a more tractable configuration of prediction in a discretized space. In fact, joint predictions do not estimate a unique location but rather a per voxel confidence, which makes it easier for a network to learn the target function [29]. In the context of 2D HPE, the benefits of predicting confidences for each pixel instead of image coordinates are well known [31, 45]. Moreover, in a multi-person environment, directly regressing the joint coordinates is unfeasible when the number of people is not known a priori, making volumetric heatmaps a natural choice for tackling bottom-up multi-person 3D HPE.

The major disadvantage of this representation is that it is memory and computational demanding, requiring some compromise during implementation that limits its full potential. Some of those compromises consist in utilizing low resolution heatmaps that introduce quantization errors or complex training strategies that involve coarse-to-fine predictions through iterative refining of network output [29].

3.2. Volumetric Heatmap Autoencoder

To overcome the aforementioned limitations without introducing quantization errors or training complexity, we propose to map volumetric heatmaps to a more tractable representation. Inspired by [17], we propose a multiple branches Volumetric Heatmap Autoencoder (VHA) that takes a set of N volumetric heatmaps \mathfrak{h} as input. At first, the volumetric

heatmaps $\{h_1, \dots, h_N\}$ are processed independently with a 2D convolutional block (*e-c2d*) in which the kernel does not move along the D dimension. In order to capture the mutual influence between joints locations, the obtained maps are then stacked along a fourth dimension and processed by a subsequent set of 3D convolutions (*e-c3d*). The resulting encoded representation, $e(\mathfrak{h})$ is finally decoded by its mirrored architecture $d(e(\mathfrak{h})) = \tilde{\mathfrak{h}}$. The general structure of the model is outlined in Fig. 2 top.

The goal of the VHA is therefore to learn a compressed representation of the input volumetric heatmaps that preserve their information content, which results in the preservation of the position of the Gaussian peaks of the various joints in the original maps. For the purpose, we maximize the F1-score, $F1(Q_{\mathfrak{h}}, Q_{\tilde{\mathfrak{h}}})$, between the set of ground truth peaks ($Q_{\mathfrak{h}}$) and the set of the decoded maps ($Q_{\tilde{\mathfrak{h}}}$). We define the set of peaks as follows:

$$Q_{\mathfrak{h}} = \bigcup_{n=1, \dots, N} \{\mathbf{u} : h_n(\mathbf{u}) > \mathbf{u}' \ \forall \mathbf{u}' \in \mathfrak{N}_{\bar{\mathbf{u}}}\} \quad (3)$$

where $\mathfrak{N}_{\bar{\mathbf{u}}}$ is the 6-connected neighborhood of $\bar{\mathbf{u}}$, i.e. the set of coordinates $\mathfrak{N}_{\bar{\mathbf{u}}} = \{\mathbf{u} : \|\mathbf{u} - \bar{\mathbf{u}}\| = 1\}$ at unit distance from $\bar{\mathbf{u}}$. Since the procedure for extracting the coordinate sets from the volumetric heatmaps is not differentiable, the former objective cannot be directly optimized as a loss component for training the VHA. To address this issue, we propose to use mean squared error (MSE) loss between \mathfrak{h} and $\tilde{\mathfrak{h}}$ as training loss.

Note that our proposed mapping purposely reduces the volumetric heatmap’s fourth dimension, making its shape coherent with the output of 2D convolutions and thus exploitable by regular CNN backbones. Additional architecture details can be found in the supplementary material.

3.3. Code Predictor and Body Joints Association

The input of the Code Predictor is represented by a RGB image, \mathbf{I} , while its output, $f(\mathbf{I})$, aims to predict the codes obtained with the VHA, Fig. 2. The architecture, Fig. 2 bottom, is inspired by [49] thus composed by a pre-trained feature extractor (convolutional part of Inception v3 [42]), and a fully convolutional block (*f-c2d*) composed of four convolutions. We trained the Code Predictor by minimizing the MSE loss between $f(\mathbf{I})$ and $e(\mathfrak{h})$, where \mathfrak{h} is the volumetric heatmap associated with the image \mathbf{I} .

At inference time, the pseudo-3D coordinates of the body joints are obtained from the decoded volumetric heatmap $\tilde{\mathfrak{h}} = d(f(\mathbf{I}))$ through a local maxima search. Eventually, if camera parameters are available, the pinhole camera equations recover the true three-dimensional coordinates of the detected joints. Additional details in the supplementary material.

model	bottleneck size	F1 on JTA			F1 on Panoptic			F1 on Human3.6m		
		@0vx	@1vx	@2vx	@0vx	@1vx	@2vx	@0vx	@1vx	@2vx
VHA ⁽¹⁾	$\frac{D}{2} \times \frac{H'}{2} \times \frac{W'}{2}$	97.1	98.4	98.5	-	-	-	-	-	-
VHA ⁽²⁾	$\frac{D}{4} \times \frac{H'}{4} \times \frac{W'}{4}$	92.5	97.0	97.1	97.1	98.6	98.9	100.0	100.0	100.0
VHA ⁽³⁾	$\frac{D}{8} \times \frac{H'}{8} \times \frac{W'}{8}$	56.5	90.3	92.9	91.9	98.7	99.6	99.7	100.0	100.0

Table 2: VHA bottleneck/code size and performances on the JTA, Panoptic and Human3.6m (protocol P2) test set in terms of F1 score at different thresholds @0, @1, and @2 voxel(s); @ t indicates that a predicted joint is considered “true positive” if the distance from the corresponding ground truth joint is less than t

As in almost all recent 2D HPE bottom-up approaches [3, 9, 5] (i.e. methods which does not require a people detection step) detected joints have to be linked together to obtain people skeletal representations. In a single person context, joint association is trivial. On the other hand, in a multi-person environment, linking joints is significantly more challenging. For the purpose, we rely on a simple distance-based heuristic where, starting from detected heads (i.e. the joint with the highest confidence), we connect the remaining $(N - 1)$ joints by selecting the closest ones in terms of 3D Euclidean distance. Associations are further refined by rejecting those that violates anatomical constraints (e.g. length of a limb greater than a certain threshold). Despite its simplicity, this approach is particularly effective when 3D coordinates of body joints are available, especially in surveillance scenarios where proxemics dynamics often regulate the spatial relationships between different individuals. Additional details are reported in the supplementary material.

3.4. Pose Refiner

The predicted 3D poses are subsequently refined by a MLP network trained to account for miss-detections and location errors. The objective of the Pose Refiner is indeed to make sure that the detected poses are complete (i.e. all the N joints are always present). To better understand how the Pose Refiner works, we define the concept of *3D poses* and *root-relative poses*. Given a person k , its 3D pose is the set $\mathbf{p}^{(k)} = \{\mathbf{x}_n^{(k)}, n = 1, \dots, N\}$ of the 3D coordinates of its N joints. The corresponding root-relative pose is then given by:

$$\mathbf{p}_{\text{rr}}^{(k)} = \left\{ \frac{\mathbf{x}_n^{(k)} - \mathbf{x}_1^{(k)}}{l_n}, n = 2, \dots, N \right\} \quad (4)$$

where \mathbf{x}_1 are the 3D coordinates of the root joint (“head-top” in our experiments) and l_n is a normalization constant computed on the training set as the maximum length of the vector that points from the root joint to any other joint of the same person.

The Pose Refiner is hence trained with MSE loss taking as input the root-relative version of the 3D poses with randomly removed joints, and an additional Gaussian noise applied to the coordinates. Given the 3D position of the root joint and the refined poses, it is straightforward to re-obtain the corresponding 3D poses by using Eq. (4).

4. Experiments

A series of experiments have been conducted on two multi-person datasets, namely JTA [8] and CMU Panoptic [12, 40, 13], as well as one well established single-person benchmark: Human3.6m [11].

JTA is a large synthetic dataset for multi-person HPE and tracking in urban scenarios. It is composed of 512 Full HD videos, 30s long, each containing an average of 20 people per frame. Due to its recent publication date, this dataset does not have a public leaderboard and it is not mentioned in other comparable HPE works. Despite this limitation, we believe it is crucial to test LoCO on JTA because it is much more complex and challenging than older benchmarks.

CMU Panoptic is another large dataset containing both single-person and multi-person sequences for a total of 65 sequences (5.5 hours of video). It is less challenging than JTA as the number of people per frame is much more limited, but it is currently the largest real-world multi-person dataset with 3D annotations.

To further demonstrate the generalization capabilities of LoCO, we also provide a direct comparison with other HPE approaches on the single person task. Without any modification to the multi-person pipeline, we achieve state of the art results on the popular Human3.6m dataset.

For each dataset we also show the upper bound obtained by using the GT volumetric heatmaps in order to highlight the strengths of this data representation. In all the following tables, we will indicate with LoCO⁽ⁿ⁾ our complete HPE pipeline, composed of the Code Predictor, the decoder of VHA⁽ⁿ⁾ and the subsequent post-processing. LoCO⁽ⁿ⁾+ is the same system with the addition of the Pose Refiner.

For all the experiments in the paper we utilized Adam optimizer with learning rate 10^{-4} . We employed batch size 1 when training the VHA and batch size 8 when training the

	PR	RE	F1	PR	RE	F1	PR	RE	F1
	@0.4 m			@0.8 m			@1.2 m		
Location Maps [21, 22]	5.80	5.33	5.42	24.06	21.65	22.29	41.43	36.96	38.26
Location Maps [21, 22] + ref. [33] + [19]	5.82	5.89	5.77	23.28	23.51	23.08	38.85	39.17	38.49
Uncompr. Volumetric Heatmaps	75.88	28.36	39.14	92.85	34.17	47.38	96.33	35.33	49.03
LoCO ⁽¹⁾	25.37	24.40	24.47	45.40	43.11	43.51	55.55	52.44	53.08
LoCO ⁽¹⁾ +	48.10	42.73	44.76	65.63	58.58	61.24	72.44	64.84	67.70
LoCO ⁽²⁾	49.37	43.45	45.73	66.87	59.02	62.02	73.54	65.07	68.29
LoCO ⁽²⁾ +	54.76	46.94	50.13	70.67	60.48	64.62	77.00	65.92	70.40
LoCO ⁽³⁾	55.37	47.84	50.82	70.63	60.94	64.76	76.81	66.31	70.44
LoCO ⁽³⁾ +	48.18	41.97	44.49	66.96	58.22	61.77	74.43	64.71	68.65
LoCO ⁽³⁾ +	49.15	42.84	45.36	67.16	58.45	61.92	74.39	64.76	68.57
GT Location Maps [21, 22]	76.07	64.83	69.59	76.07	64.83	69.59	76.07	64.83	69.59
GT Volumetric Heatmaps	99.96	99.96	99.96	99.99	99.99	99.99	99.99	99.99	99.99

Table 3: Comparison of our LoCO approach with other strong baselines and competitors on the JTA test set. In PR (precision), RE (recall) and F1, @ t indicates that a predicted joint is considered “true positive” if the distance from the corresponding ground truth joint is less than t . Last two rows contain the upper bounds obtained using the ground truth location maps and volumetric heatmaps respectively

Code Predictor. We employed Inception v3 [42] as backbone for the Code Predictor, which is followed by 3 convolutions with ReLU activation having kernel size 4 and with 1024, 512 and 256 channels respectively. A last 1×1 convolution is performed to match the compressed volumetric heatmap’s number of channels. Additional training details in the supplementary material.

4.1. Compression Levels

In order to understand how different code sizes in the VHA affects the performance of our Code Predictor network, multiple VHA versions have been tested. Specifically, we designed three VHA versions with decreasing bottleneck sizes. Each version has been trained on JTA first and then finetuned on CMU Panoptic and Human3.6m. VHA’s architecture details are depicted in Tab. 1 for every version.

As shown in Tab. 2, as the bottleneck size decreases, there is a corresponding decrease in the F1-score. Intuitively, the more we compress, the less information is being preserved. VHA⁽¹⁾ is only considered when using JTA, as VHA⁽²⁾ and VHA⁽³⁾ already obtain an almost lossless compression on Panoptic and Human3.6m, due to their smaller number of people in the scene.

All the experiments has been conducted considering a 14 joints volumetric heatmap representation of shape $14 \times D \times H' \times W'$, where H' and W' are height and width downsampled by a factor of 8, while D has been fixed to 316 bins. Note that the real-world depth grid covered by our representation is a uniform discretization in $[0, 100]$ m for JTA, $[0, 7]$ m for Panoptic and $[1.8, 8.1]$ m for Human3.6m. Thus, every bin has a depth size of approximately 0.32m for JTA and 0.02m for Panoptic and Human3.6m.

4.2. HPE Experiments on JTA Dataset

On the JTA dataset we compared LoCO against the Location Maps based approaches of [21, 22]. Currently the Location Maps representation is the most relevant alternative to volumetric heatmaps to approach the 3D HPE task in a bottom-up fashion and therefore represents our main competitor.

A Location Maps is a per-joint feature channel that stores the 3D coordinate x , y , or z at the joint 2D pixel location. For each joint there are three location-maps and the 2D heatmap. The 2D heatmap encodes the pixel location of the joint as a confidence map in the image plane. The 3D position of a joint can then be obtained from its Location Map at the 2D pixel location of the joint. For a fair comparison, we utilized the same network (Inception v3 + f -c2d) to directly predict the Location Maps. The very low F1 score demonstrate that Location Maps are not suitable for images with multiple overlapping people, not being able to effectively handle the challenging situations peculiar of crowded surveillance scenarios (see Tab. 3).

Additionally, we report a comparison with a strong top-down baseline that uses YOLOv3 [33] for the people detection part and [19] as the single-person pose estimator. [19], like almost all single person methods, provides root-relative joint coordinates and not the absolute 3D position. We thus performed the 3D alignment according to [37] by minimizing the distance between 2D pose and re projected 3D pose. We outperform this top-down pipeline by a large margin in terms of F1-score, while being significantly faster; LoCO is able to process Full HD images with more than 50 people at 8 FPS on a Tesla V100 GPU, while the top-down baseline runs at an average of 0.5 FPS (16 times slower). The re-

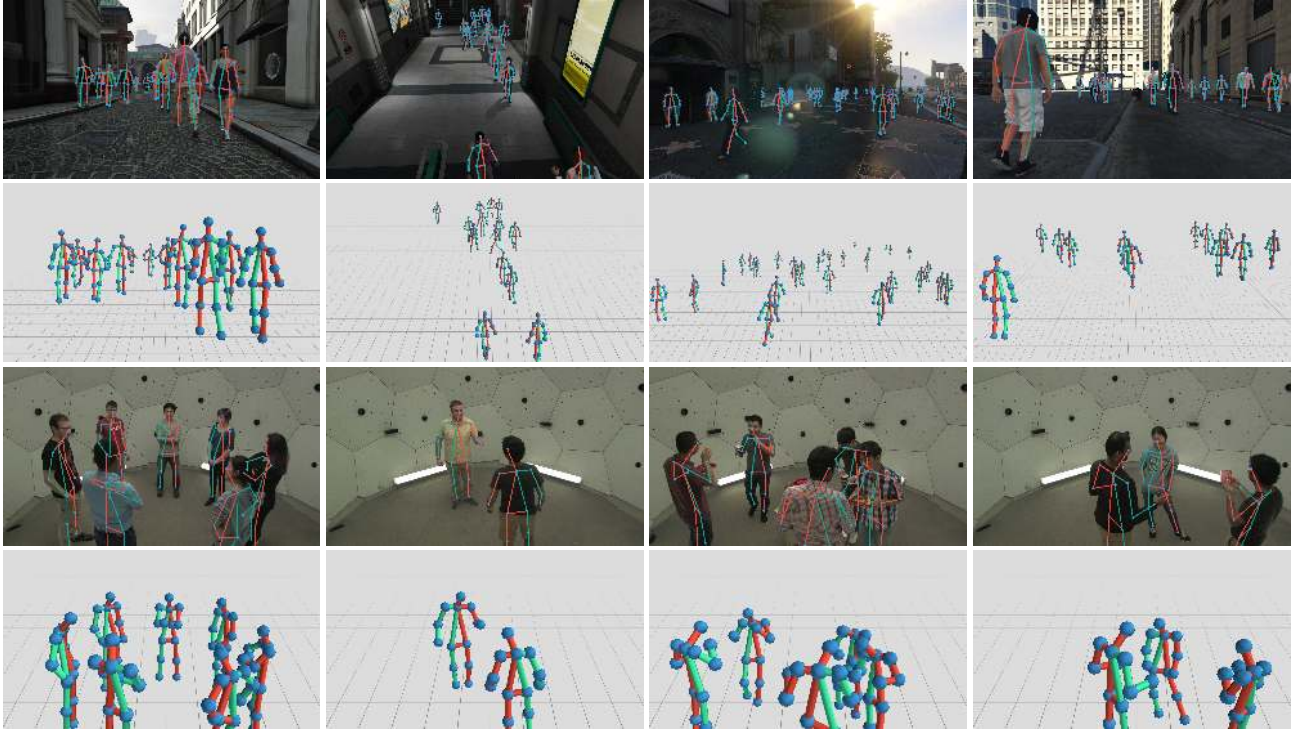


Figure 3: Qualitative results of LoCO⁽²⁾+ on the JTA and Panoptic datasets. We show both the 3D poses (JTA: 2nd row, Panoptic: 4th row) and the corresponding 2D versions re-projected on the image plane (JTA: 1st row, Panoptic: 3rd row)

call gap is mostly due to the fact that the detection phase in top-down approaches usually miss overlapped or partially occluded people on the crowded JTA scenes.

Finally, we compared against an end-to-end model trained to directly predict the volumetric heatmaps without compression (“Uncompr. Volumetric Heatmaps” in Tab. 3). Specifically, we stacked the Code Predictor and the VHA⁽²⁾’s decoder and trained it in an end-to-end fashion. Our technique outperforms this version at every compression rate. In fact, the sparseness of the target makes it difficult to effectively exploit the redundancy of body poses in the ground truth annotation leading to a more complex training phase.

We point out that LoCO⁽²⁾+ obtains by far the best result in terms of F1-score compared to all evaluated approaches and baselines, thus demonstrating the effectiveness of our method. Moreover, the best result has been obtained using the VHA⁽²⁾’s mapping, which seemingly exhibits the best compromise between information preserved and density of representation. It is also very interesting to note that the upper bound for Volumetric Heatmaps is much higher than that of Location Maps (last two rows of Tab. 3), highlighting the superiority of volumetric heatmaps in crowded scenarios. It is finally worth noticing that LoCO⁽¹⁾+ and LoCO⁽³⁾+ obtain very close results, indicating that an ex-

tremely lossy compression can lead to a poor solution as much as utilizing a too sparse and oversized representation.

Following the protocol in [8], we trained all our models (and those with Location Maps) on the 256 sequences of the JTA training set and tested our complete pipeline only on every 10th frame of the 128 test sequences. Qualitative results are presented in Fig. 3.

4.3. HPE Experiments on Panoptic Dataset

Here we propose a comparison between LoCO and three strong multi-person approaches [53, 52, 32] on CMU Panoptic following the test protocol defined in [52]. The results, shown in the Tab. 4, are divided by action type and are expressed in terms of Mean Per Joint Position Error (MPJPE). MPJPE is calculated by firstly associating predicted and ground truth poses, by means of a simple Hungarian algorithm. In the Tab. 4 we also report the F1-score: the solely MPJPE metric is not meaningful as it does not take into account missing detections or false positive predictions.

The obtained results show the advantages of using volumetric heatmaps for 3D HPE, as LoCO⁽²⁾+ achieves the best result in terms of average MPJPE on the Panoptic test set. For the sake of fairness, we also tested on the no longer maintained “mafia” sequence. However, the older version

	MPJPE [mm]					F1
	Haggl.	Mafia	Ultim.	Pizza	Mean	
[32]	218	187	194	221	203	-
[52]	140	166	151	156	153	-
[53]	72	79	67	94	72	-
LoCO ⁽²⁾ +	45	95	58	79	69	89.21
LoCO ⁽³⁾ +	48	105	63	91	77	87.87
GT	9	12	9	9	10	100

Table 4: Comparison on the CMU Panoptic dataset. Results are shown in terms of MPJPE [mm] and F1 detection score. Last row: results with ground truth volumetric heatmaps

	method	N	P1	P1 (a)	P2	P2 (a)
top-down	Rogez <i>et al.</i> [36]	13	63.2	53.4	87.7	71.6
	Dabral <i>et al.</i> [6]	16	-	-	-	65.2
	Rogez <i>et al.</i> [37]	13	54.6	45.8	65.4	54.3
	Moon <i>et al.</i> [23]	17	35.2	34.0	54.4	53.3
bottom-up	Mehta <i>et al.</i> [22]	17	-	-	80.5	-
	Mehta <i>et al.</i> [21]	17	-	-	69.9	-
	LoCO ⁽²⁾ +	14	84.0	75.4	96.6	77.1
	LoCO ⁽³⁾ +	14	51.1	43.4	61.0	49.1
	GT Vol. Heatmaps	14	15.6	14.9	15.0	14.3

Table 5: Comparison on the Human3.6m dataset in terms of average MPJPE [mm]. “(a)” indicates the addition of rigid alignment to the test protocol; N is the number of joints considered by the method. Last row: results with ground truth volumetric heatmaps

of the dataset utilizes a different convention for the joint positions. This, in fact, is reflected by the worst performance in that sequence only. Once again, the best trade-off is obtained using VHA⁽²⁾, due to VHA⁽³⁾’s mapping partial loss of information. The GT upper bound in Tab. 4 further demonstrate the potential of our representation. Qualitative results are presented in Fig. 3.

4.4. HPE Experiments on Human3.6m Dataset

In analogy with previous experiments, we tested LoCO on Human3.6m. Unlike most existing approaches, we apply our multi-person method as it is, without exploiting the knowledge of the single-person nature of the dataset, as we want to demonstrate its effectiveness even in this simpler context. Results, with and without rigid alignment, are reported in terms of MPJPE following the P1 and P2 protocols. In the P1 protocol, six subjects (S1, S5, S6, S7, S8 and S9) are used for training and every 64th frame of subject S11/camera 2 is used for testing. For the P2 protocol, all the frames from subjects S9 and S11 are used for testing and only S1, S5, S6, S7 and S8 are used for training.

Tab. 5 shows a comparison with recent state-of-the-art multi-person methods, showing that our method is well

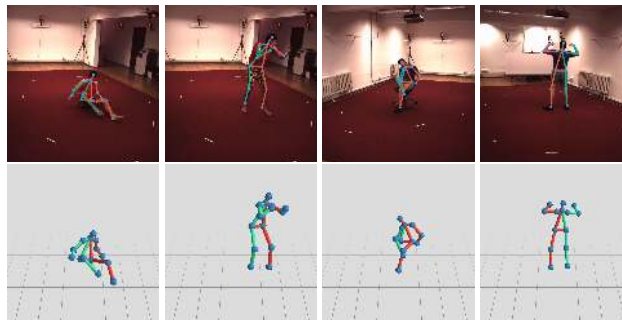


Figure 4: Qualitative results of LoCO⁽³⁾+

suites even in the single person context, as LoCO⁽³⁾+

5. Discussion and Conclusions

In conclusion, we presented a single-shot bottom-up approach for multi-person 3D HPE suitable for both crowded surveillance scenarios and for simpler, even single person, contexts without any changes. Our LoCO approach allows us to exploit volumetric heatmaps as a ground truth representation for the 3D HPE task. Instead, without compression, this would lead to a sparse and extremely high dimensional output space with consequences on both the network size and the stability of the training procedure. In comparison with top-down approaches, we removed the dependency on the people detector stage, hence gaining both in terms of robustness and assuring a constant processing time at the increasing of people in the scene. The experiments show state-of-the-art performance on all the considered datasets. We also believe that this new simple compression strategy can foster future research by enabling the full potential of the volumetric heatmap representation in contexts where it was previously intractable.

Acknowledgments

This work was supported by Panasonic Corporation and by the Italian Ministry of Education, Universities and Research under the project COSMOS PRIN 2015 programme 201548C5NT.

References

- [1] H. Arai, Y. Chayama, H. Iyatomi, and K. Oishi. Significant dimension reduction of 3d brain mri using 3d convolutional autoencoders. In *2018 40th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pages 5162–5165, July 2018. 1
- [2] Federica Bogo, Angjoo Kanazawa, Christoph Lassner, Peter Gehler, Javier Romero, and Michael J Black. Keep it smpl: Automatic estimation of 3d human pose and shape from a single image. In *European Conference on Computer Vision (ECCV)*, 2016. 2
- [3] Zhe Cao, Tomas Simon, Shih-En Wei, and Yaser Sheikh. Realtime multi-person 2d pose estimation using part affinity fields. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 1, 5
- [4] Ching-Hang Chen and Deva Ramanan. 3d human pose estimation=2d pose estimation+ matching. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 2
- [5] Yilun Chen, Zhicheng Wang, Yuxiang Peng, Zhiqiang Zhang, Gang Yu, and Jian Sun. Cascaded pyramid network for multi-person pose estimation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 5
- [6] Rishabh Dabral, Anurag Mundhada, Uday Kusupati, Safeer Afaque, Abhishek Sharma, and Arjun Jain. Learning 3d human pose from structure and motion. In *European Conference on Computer Vision (ECCV)*, 2018. 2, 8
- [7] Angela Dai and Matthias Nießner. 3dmv: Joint 3d-multi-view prediction for 3d semantic scene segmentation. In *European Conference on Computer Vision (ECCV)*, 2018. 3
- [8] Matteo Fabbri, Fabio Lanzi, Simone Calderara, Andrea Palazzi, Roberto Vezzani, and Rita Cucchiara. Learning to detect and track visible and occluded body joints in a virtual world. In *European Conference on Computer Vision (ECCV)*, 2018. 2, 5, 7
- [9] Mihai Fieraru, Anna Khoreva, Leonid Pishchulin, and Bernt Schiele. Learning to refine human pose estimation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2018. 5
- [10] Mir Rayat Intiaz Hossain and James J Little. Exploiting temporal information for 3d human pose estimation. In *European Conference on Computer Vision (ECCV)*, 2018. 2
- [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*, 2014. 1, 2, 5
- [12] Hanbyul Joo, Hao Liu, Lei Tan, Lin Gui, Bart Nabbe, Iain Matthews, Takeo Kanade, Shohei Nobuhara, and Yaser Sheikh. Panoptic studio: A massively multiview system for social motion capture. In *The IEEE International Conference on Computer Vision (ICCV)*, 2015. 2, 5
- [13] Hanbyul Joo, Tomas Simon, Xulong Li, Hao Liu, Lei Tan, Lin Gui, Sean Banerjee, Timothy Scott Godisart, Bart Nabbe, Iain Matthews, Takeo Kanade, Shohei Nobuhara, and Yaser Sheikh. Panoptic studio: A massively multiview system for social interaction capture. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2017. 5
- [14] Angjoo Kanazawa, Michael J Black, David W Jacobs, and Jitendra Malik. End-to-end recovery of human shape and pose. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 2
- [15] Kyoungoh Lee, Inwoong Lee, and Sanghoon Lee. Propagating lstm: 3d pose estimation based on joint interdependency. In *European Conference on Computer Vision (ECCV)*, 2018. 2
- [16] Mude Lin, Liang Lin, Xiaodan Liang, Keze Wang, and Hui Cheng. Recurrent 3d pose sequence machines. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 2
- [17] Wenjie Luo, Bin Yang, and Raquel Urtasun. Fast and furious: Real time end-to-end 3d detection, tracking and motion forecasting with a single convolutional net. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3569–3577, 2018. 4
- [18] Diogo C Luvizon, David Picard, and Hedi Tabia. 2d/3d pose estimation and action recognition using multitask deep learning. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 2
- [19] Julieta Martinez, Rayat Hossain, Javier Romero, and James J Little. A simple yet effective baseline for 3d human pose estimation. In *International Conference on Computer Vision (ICCV)*, 2017. 2, 6
- [20] Dushyant Mehta, Helge Rhodin, Dan Casas, Pascal Fua, Oleksandr Sotnychenko, Weipeng Xu, and Christian Theobalt. Monocular 3d human pose estimation in the wild using improved cnn supervision. In *International Conference on 3D Vision (3DV)*, 2017. 2
- [21] Dushyant Mehta, Oleksandr Sotnychenko, Franziska Mueller, Weipeng Xu, Srinath Sridhar, Gerard Pons-Moll, and Christian Theobalt. Single-shot multi-person 3d pose estimation from monocular rgb. In *International Conference on 3D Vision (3DV)*, 2018. 2, 3, 6, 8
- [22] Dushyant Mehta, Srinath Sridhar, Oleksandr Sotnychenko, Helge Rhodin, Mohammad Shafiei, Hans-Peter Seidel, Weipeng Xu, Dan Casas, and Christian Theobalt. Vnect: Real-time 3d human pose estimation with a single rgb camera. *ACM Transactions on Graphics (TOG)*, 2017. 2, 6, 8
- [23] Gyeongsik Moon, Ju Yong Chang, and Kyoung Mu Lee. Camera distance-aware top-down approach for 3d multi-person pose estimation from a single rgb image. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 10133–10142, 2019. 2, 8
- [24] Francesc Moreno-Noguer. 3d human pose estimation from a single image via distance matrix regression. In *The IEEE Conference on Computer Vision and Pattern Recognition*, 2017. 2
- [25] Aiden Nibali, Zhen He, Stuart Morgan, and Luke Prendergast. 3d human pose estimation with 2d marginal heatmaps. *IEEE Winter Conference on Applications of Computer Vision (WACV)*, 2019. 2
- [26] Bruce Xiaohan Nie, Ping Wei, and Song-Chun Zhu. Monocular 3d human pose estimation by predicting depth on joints.

- In *International Conference on Computer Vision (ICCV)*, 2017. 2
- [27] Mohamed Omran, Christoph Lassner, Gerard Pons-Moll, Peter Gehler, and Bernt Schiele. Neural body fitting: Unifying deep learning and model based human pose and shape estimation. In *International Conference on 3D Vision (3DV)*, 2018. 2
- [28] Georgios Pavlakos, Xiaowei Zhou, and Kostas Daniilidis. Ordinal depth supervision for 3d human pose estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7307–7316, 2018. 2
- [29] Georgios Pavlakos, Xiaowei Zhou, Konstantinos G Derpanis, and Kostas Daniilidis. Coarse-to-fine volumetric prediction for single-image 3d human pose. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 1, 2, 3, 4
- [30] Georgios Pavlakos, Luyang Zhu, Xiaowei Zhou, and Kostas Daniilidis. Learning to estimate 3d human pose and shape from a single color image. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 2
- [31] Tomas Pfister, James Charles, and Andrew Zisserman. Flowing convnets for human pose estimation in videos. In *Proceedings of the IEEE International Conference on Computer Vision*, 2015. 4
- [32] Alin-Ionut Popa, Mihai Zanfir, and Cristian Sminchisescu. Deep multitask architecture for integrated 2d and 3d human sensing. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017. 7, 8
- [33] Joseph Redmon and Ali Farhadi. Yolov3: An incremental improvement. *arXiv preprint arXiv:1804.02767*, 2018. 6
- [34] Helge Rhodin, Mathieu Salzmann, and Pascal Fua. Unsupervised geometry-aware representation for 3d human pose estimation. *European Conference on Computer Vision (ECCV)*, 2018. 2
- [35] Helge Rhodin, Jörg Spörri, Isinsu Katircioglu, Victor Constantin, Frédéric Meyer, Erich Müller, Mathieu Salzmann, and Pascal Fua. Learning monocular 3d human pose estimation from multi-view images. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 2
- [36] Gregory Rogez, Philippe Weinzaepfel, and Cordelia Schmid. Lcr-net: Localization-classification-regression for human pose. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 8
- [37] Gregory Rogez, Philippe Weinzaepfel, and Cordelia Schmid. Lcr-net++: Multi-person 2d and 3d pose detection in natural images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2019. 2, 3, 6, 8
- [38] István Sáráandi, Timm Linder, Kai O Arras, and Bastian Leibe. Synthetic occlusion augmentation with volumetric heatmaps for the 2018 eccv posetrack challenge on 3d human pose estimation. *European Conference on Computer Vision (ECCV) - Workshops*, 2018. 1, 2
- [39] Matthias Scholz, Martin Fraunholz, and Joachim Selbig. Nonlinear principal component analysis: neural network models and applications. In *Principal Manifolds for Data Visualization and Dimension Reduction*, 2008. 1
- [40] Tomas Simon, Hanbyul Joo, and Yaser Sheikh. Hand key-point detection in single images using multiview bootstrapping. *CVPR*, 2017. 5
- [41] Xiao Sun, Bin Xiao, Shuang Liang, and Yichen Wei. Integral human pose regression. *European Conference on Computer Vision (ECCV)*, 2018. 1, 2
- [42] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 4, 6
- [43] Wei Tang, Pei Yu, and Ying Wu. Deeply learned compositional models for human pose estimation. In *European Conference on Computer Vision (ECCV)*, 2018. 1
- [44] Bugra Tekin, Pablo Marquez Neila, Mathieu Salzmann, and Pascal Fua. Learning to fuse 2d and 3d image cues for monocular body pose estimation. In *International Conference on Computer Vision (ICCV)*, 2017. 2
- [45] Jonathan J Tompson, Arjun Jain, Yann LeCun, and Christoph Bregler. Joint training of a convolutional network and a graphical model for human pose estimation. In *Advances in neural information processing systems*, 2014. 4
- [46] Róbert Torfason, Fabian Mentzer, Eiríkur Ágústsson, Michael Tschannen, Radu Timofte, and Luc Van Gool. Towards image understanding from deep compression without decoding. In *International Conference on Learning Representations*, 2018. 2
- [47] Jing Wang, Haibo He, and Danil V Prokhorov. A folded neural network autoencoder for dimensionality reduction. *Procedia Computer Science*, 2012. 1
- [48] Tongzhou Wang, Jun-Yan Zhu, Antonio Torralba, and Alexei A Efros. Dataset distillation. *arXiv preprint arXiv:1811.10959*, 2018. 2
- [49] Bin Xiao, Haiping Wu, and Yichen Wei. Simple baselines for human pose estimation and tracking. In *European Conference on Computer Vision (ECCV)*, 2018. 1, 4
- [50] Wei Yang, Wanli Ouyang, Xiaolong Wang, Jimmy Ren, Hongsheng Li, and Xiaogang Wang. 3d human pose estimation in the wild by adversarial learning. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2018. 2
- [51] Hashim Yasin, Umar Iqbal, Bjorn Kruger, Andreas Weber, and Juergen Gall. A dual-source approach for 3d pose estimation from a single image. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 2
- [52] Andrei Zanfir, Elisabeta Marinoiu, and Cristian Sminchisescu. Monocular 3d pose and shape estimation of multiple people in natural scenes - the importance of multiple scene constraints. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 2, 3, 7, 8
- [53] Andrei Zanfir, Elisabeta Marinoiu, Mihai Zanfir, Alin-Ionut Popa, and Cristian Sminchisescu. Deep network for the integrated 3d sensing of multiple people in natural images. In *Advances in Neural Information Processing Systems*, 2018. 2, 3, 7, 8
- [54] Xingyi Zhou, Qixing Huang, Xiao Sun, Xiangyang Xue, and Yichen Wei. Towards 3d human pose estimation in the wild: a weakly-supervised approach. In *International Conference on Computer Vision (ICCV)*, 2017. 2