



A generalizable approach for multi-view 3D human pose regression

Abdolrahim Kadkhodamohammadi¹ · Nicolas Padoy¹

Received: 13 March 2019 / Revised: 8 April 2020 / Accepted: 24 August 2020 / Published online: 8 October 2020
© Springer-Verlag GmbH Germany, part of Springer Nature 2020

Abstract

Despite the significant improvement in the performance of monocular pose estimation approaches and their ability to generalize to unseen environments, multi-view approaches are often lagging behind in terms of accuracy and are specific to certain datasets. This is mainly due to the fact that (1) contrary to real-world single-view datasets, multi-view datasets are often captured in controlled environments to collect precise 3D annotations, which do not cover all real-world challenges, and (2) the model parameters are learned for specific camera setups. To alleviate these problems, we propose a two-stage approach to detect and estimate 3D human poses, which separates single-view pose detection from multi-view 3D pose estimation. This separation enables us to utilize each dataset for the right task, i.e. single-view datasets for constructing robust pose detection models and multi-view datasets for constructing precise multi-view 3D regression models. In addition, our 3D regression approach only requires 3D pose data and its projections to the views for building the model, hence removing the need for collecting annotated data from the test setup. Our approach can therefore be easily generalized to a new environment by simply projecting 3D poses into 2D during training according to the camera setup used at test time. As 2D poses are collected at test time using a single-view pose detector, which might generate inaccurate detections, we model its characteristics and incorporate this information during training. We demonstrate that incorporating the detector's characteristics is important to build a robust 3D regression model and that the resulting regression model generalizes well to new multi-view environments. Our evaluation results show that our approach achieves competitive results on the Human3.6M dataset and significantly improves results on a multi-view clinical dataset that is the first multi-view dataset generated from live surgery recordings.

Keywords Multi-view human pose estimation · 3D pose regression · Neural networks · Generalizability

1 Introduction

Single-view human detection and body pose estimation have enjoyed a great deal of attention over the last decades in the field of computer vision because of their importance for various applications, ranging from activity recognition to human computer interaction. More recently, the emergence of deep learning has pushed the boundaries in many fields, including computer vision. The combination of deep learning with the availability of large datasets, such as *MPII Pose* [4] and *MS COCO* [27], has spawned many promising approaches for single-view human detection and pose estimation [10,35,50]. But the presence of clutter and occlusions degrades their per-

formance. Capturing an environment from complementary views permits to reduce the risk of occlusions, especially in busy environments, as shown in Fig. 1. In addition, the availability of calibrated multi-view data greatly facilitates the process of lifting 2D scenes into 3D, which is important for many applications such as augmented reality.

Despite the inherent benefits of capturing an environment from multiple views, multi-view approaches have not achieved the same level of maturity as compared to single-view approaches, mostly due to two reasons: firstly, multi-view datasets are generally recorded in controlled environments in order to use motion capture systems to acquire precise 3D ground truth location data. This removes the need for the tedious and error-prone manual annotation of the abundant number of frames coming from all views for generating ground truth 3D poses. Even though there are large multi-view datasets such as *Human3.6M* [23] and *HumanEva* [40], the simple backgrounds and tight clothes required by motion capture systems make these datasets

✉ Abdolrahim Kadkhodamohammadi
kadkhodamohammad@unistra.fr

Nicolas Padoy
npadoy@unistra.fr

¹ ICube, University of Strasbourg, CNRS, IHU Strasbourg, Strasbourg, France




Fig. 1 A set of images captured by a multi-view camera system at the same time step. Even though some body parts are occluded in one view due to self- or object-occlusion, they might appear in other complementary views

trivial for 2D pose estimation methods. Monocular pose estimation approaches report low 2D body part localization errors even without finetuning [11,31]. For these reasons, single- and multi-view pose estimation models trained on datasets captured in such controlled laboratory environments do not generalize well to real-world data, which is often visually much more complex due to occlusions, clutter and the presence of multiple persons in the scene. Secondly, current multi-view approaches [13,37,41] learn model parameters that are specific to each multi-view camera setup. In other words, to apply these approaches on a *new* multi-view scenario, it is required to collect new annotated data that includes both multi-view images and their corresponding 3D ground truth poses for the same camera setup. On the one hand, generating synthetic datasets for these approaches would require not only the generation of 3D body poses, but also of photo-realistic rendering of humans with different shapes, textures and backgrounds to allow generalization to the real world, which is not a trivial task. On the other hand, generating such training data using either motion capture systems or manual annotations, especially in the case of data-hungry deep learning methods, is not always feasible in uncontrolled environments and very tedious. We therefore propose an approach that benefits from existing multi-view datasets to perform multi-view 3D pose estimation in new multi-view setups.

Our approach formulates the problem of multi-view 3D pose estimation in a two-step framework: (1) single-view pose detections and (2) multi-view 3D pose regression. We separate these two steps for two reasons. First, we can better exploit available single-view and multi-view datasets for the right task. Single-view datasets, such as *MPII Pose* [4] and *MS COCO* [27], include diverse and challenging frames from everyday activities or movies originating from amateur to professional recordings. Therefore, models trained on these datasets can better cope with real-world challenges and generalize to new environments. But, these single-view datasets are lacking 3D annotations, contrary to multi-view datasets, which often come with accurate 3D body poses. As these are, however, much simpler for the task of 2D pose esti-

mation [11,31], researchers have proposed methods to jointly use both single- and multi-view datasets in order to construct more robust 3D pose estimation models from multiple views [3,7]. Changes in camera setups, however, require the retraining of the model on training data from the *same* camera setup. This strictly limits the deployment of the models to environments where such training data exists. The second reason for our two steps approach is that we can better generalize to new multi-view environments by assuming that lifting 2D body poses into 3D is independent of the images given the 2D pose detections. This assumption implies that we do not need to collect 2D image data for training the 3D regression function and that any set of plausible 3D body poses can be used instead by computing body pose projections into 2D.

To learn a multi-view 3D regression function, we propose a method that relies on a multi-stage neural network. The input of this network is a set of corresponding multi-view 2D detections for each individual person. At test time, they are collected using a state-of-the-art single-view detector. We assume that the camera system is fully calibrated and can therefore use epipolar geometry to establish the multi-view correspondences per person. This process also allows us to detect the number of persons per multi-view frame.¹ This is in contrast to current multi-view RGB approaches, which tackle either single-person scenarios [18,20] or multi-person scenarios where the number of persons is known *a priori* [6,28].

The proposed network consists of a series of blocks of fully connected layers with intermediate supervision at each block. The input to each block is the raw network input, i.e. the concatenated 2D poses, and the output from previous block if it exists. The network can therefore build a high dimensional function and refine the output of the previous block to achieve a more reliable regression function. In order to generalize to new multi-view setups, we do not use images during training but construct training data solely by project-

¹ We define a multi-view frame as the set of all images captured from all views at the same time step.

ing Human3.6M's 3D poses. We use Human3.6M because it is the largest publicly available multi-view dataset and it includes men and women of different sizes. The projected 2D poses are generated according to the camera parameters used at test time. In practice, 2D poses are detected at test time using a 2D pose detector that may be noisy and inaccurate. In order to cope with these inaccuracies, we propose to perturb during training the 2D locations of the body joints by random noise that is generated based on the characteristics of the 2D detector. We also propose to incorporate a detection confidence for each body joint, computed based on the amount of noise added during training. This provides a representation for the detection confidence generated by the detector at test time. Therefore, the approach can take into account not only joint locations, but also detection precision to build a robust regression function.

We use two datasets to perform quantitative and qualitative evaluations and compare with state-of-the-art results on these datasets. We first report results on the Human3.6M dataset [23] to characterize the properties and the performance of our approach. This dataset includes recordings of several actions performed by professional actors of different genders. This dataset has been recorded by a fully calibrated four-view camera system and a motion capture system to collect ground truth 3D positions of the body joints. We also evaluate our approach on a challenging multi-view dataset [26,43] to show the generalization ability of our approach. This dataset is generated from real surgery recordings obtained in an operating room (OR) using a three-view camera system and hence is called *Multi-view OR* (MVOR) in the following.² Our approach improves 3D body part localization on Human3.6M and significantly reduces the localization error on the multi-view OR dataset without using any training data from this dataset.

The main contributions of the paper are twofold. First, we present a simple and yet accurate multi-view 3D pose estimation approach that can generalize well to new multi-view environments. In contrast to current state-of-the-art methods, the approach exploits an existing multi-view dataset to build models for new multi-view environments without any need for new annotation. Second, this is the first multi-view RGB approach that has been quantitatively evaluated on data captured in an unconstrained environment. Related work is discussed in Sect. 2. The proposed approach is described in Sect. 3. Section 4 presents experimental results and discussion. Finally, we conclude in Sect. 5. The list of abbreviations used in paper is presented in Table 1.

² The MVOR dataset is publicly available at <http://camma.u-strasbg.fr/datasets>.

Table 1 List of abbreviations

FC:	Fully Connected Layer
GT:	Ground Truth
MPII:	MPII Human Pose Dataset [4]
MS COCO:	Microsoft Common Objects in Context Dataset [27]
MV:	Multi-view
MVOR:	Multi-view Operating Room Dataset [43]
OR:	Operating Room
RGB:	Red Green Blue

2 Related work

Multi-view segmentation-based 3D pose estimation. Hofmann and Gavrilu [20] use foreground segmentation to estimate body silhouettes per view. Then, 3D pose candidates are obtained by matching a library of exemplars. Texture information and shape similarity across all views combined with temporal information are used to compute the final 3D poses. Similarly, Gall et al. [18] propose a two-layer framework that iteratively improves foreground segmentation and retrieved body poses by incorporating both multi-view and temporal information. Other approaches have deployed optical flow estimation [12], 2D as well as 3D motion cues [45] and low-rank multi-view feature fusion combined with sparse spectral embedding [51] to estimate 3D poses. In contrast to our work, these approaches are only evaluated on single-person datasets. More importantly, it is not always possible to compute foreground in cluttered environments, such as in operating rooms. Therefore, these approaches can only be evaluated on data recorded in environments with simple backgrounds.

Multi-view part-based 3D pose estimation Several multi-view 3D pose estimation approaches [2,3,5,7,9,26] have been proposed that rely on a part-based framework [16]. This part-based framework provides an elegant formalism to optimize over different potential functions for incorporating image features, multi-view cues, temporal information and body physical constraints. Burenus et al. [9] propose an approach that extends pictorial structures [15,17] to multi-view and to perform exact 3D inference by using simple binary pairwise potential functions. Instead, Amin et al. [2,3] use 2D inference with more complex pairwise potentials, multi-view cues and triangulation to estimate 3D poses. Huang et al. [21] have proposed to fuse both multi-view images and body-worn inertia measurement unit data to estimate 3D body poses using deep neural networks. In [42], multi-view color and depth videos are used to estimate 3D poses. Belagianis et al. [5] have also deployed different pairwise potentials for incorporating both body physical constraints and multi-view features. This approach allows to perform approximate

3D inference by selecting a limited number of hypotheses per individual. This approach has further been extended to incorporate temporal information [6] and to use a deep neural network-based body part detector [7]. Recently, Pavlakos et al. [37] have used deep neural network to predict body part score maps across all views and then estimated body poses by using a 3D pictorial structures approach.

In contrast to our work, all these approaches have only been evaluated on datasets recorded in constrained laboratory environments and also require the number of person to be known *a priori*. MVDeep3DPS presented in [26] is an exception, but this approach relies on multi-view RGB-D input to estimate 3D body poses. Additionally, all these approaches need in general to learn model parameters on data from the same camera setup. Moreover, optimizing these energy functions is demanding, especially in 3D, which makes these approaches not suitable for real-time applications. In our work, we do not require images with pose annotations from the camera setup used at test time and learn model parameters by using existing datasets. Furthermore, our approach performs both human detection and pose estimation. As our regression function uses a multi-layer neural network, it runs in super real-time on a single consumer GPU card.

Single-view 3D pose estimation. Recently, many deep learning-based approaches have been proposed to directly regress body poses in 3D from a monocular image or an image sequence. Pavlakos et al. [36] use a stack of a fully convolutional network [34] to iteratively compute 3D heatmaps per body parts. Tekin et al. [47] propose to learn an auto-encoder that maps 3D body joints into a high-dimension latent space for discovering joint dependencies and then to learn a convolutional network that maps an image into this high-dimensional pose space. In [46], motion compensation is used to align several consecutive frames and construct a rectified spatiotemporal volume that is then fed into a 3D regression function. Most recently, Luvizon et al. [29] tackled both human pose estimation and activity recognition jointly over video clips. Other approaches have built deep pose grammar representations [14], sparse representation [25], skeleton map [49] and multitask objectives [30,39] to enforce more constraints and obtain a more accurate 3D regression function. These approaches are trained on images with accurate 3D ground truth poses. The main issue is that to generate such accurate 3D annotations, motion capture systems are used in controlled laboratory environments with simple backgrounds. Models trained on such image data do not generalize well to real-world scenes.

Another line of work relies on two-stage methods, where 2D body parts are first predicted using 2D pose detectors [10,34,50] and then 3D body part locations are computed by relying on these predictions [11,31,32]. In comparison with direct 3D regression approaches, these approaches ben-

efit from the diverse, challenging and real-world datasets, e.g. MS COCO [27] and MPII Pose [4], to train reliable 2D pose detector models that generalize well. To compute 3D body locations, exemplar-based approaches are used by matching lower and upper body parts separately [24] and by matching the whole skeleton [11]. More recently, [32] proposed to regress from 2D Euclidean distance matrices (EDM) to 3D EDM instead of using traditional 2D-to-3D regression in the Cartesian coordinate system [23,38]. The regression is performed using a fully convolutional network and 3D poses are recovered via a multidimensional scaling algorithm [8]. Martinez et al. [31] showed that a simple fully connected network to regress from 2D to 3D outperforms [32] and achieves state-of-the-art results on Human3.6M. We also adopt a two-stage framework in our multi-view approach and use a fully connected network as a 2D-to-3D regression function. The single-view model in [31] was, however, trained on the output of the 2D detector used during test time. In contrast, our approach relies solely on ground truth during training and instead generates training samples that comply with the behavior of the 2D detector used at test time. This is an interesting property of our approach, which enables us to train our network on Human3.6M and test on a completely different multi-view dataset.

3 Methodology

In this section, we present our proposed approach for multi-view 3D pose estimation. We assume that we have a calibrated multi-view system recording an environment from a set of complementary views. Our objective is to detect and predict human body poses in 3D given images captured from all views. In a probabilistic formulation, we want to compute $p(Y, \mathbb{X}, \mathbb{I})$, the joint distribution over the following three random variables: (1) the 3D body poses $Y = (y_1, y_2, \dots, y_P)$, where P is the number of body joints and $y_i \in \mathbb{R}^3$ is a body joint location in 3D; (2) the 2D body poses $\mathbb{X} = (X_1, X_2, \dots, X_V)$, where V is the number of viewpoints and X_j is the tuple of pixel coordinates indicating the body joints of a 2D pose in view j ; and (3) all 2D images $\mathbb{I} = (I_1, I_2, \dots, I_V)$, where I_j is the image taken from the j^{th} viewpoint. Such a formulation makes no limiting assumption and indicates that a 3D body pose is jointly dependent on its appearance in all individual views. However, learning such a model requires collecting training data from the same multi-view setup that we want to apply the model to.

Without loss of generality, we can rewrite the joint probability distribution as:

$$p(Y, \mathbb{X}, \mathbb{I}) = p(Y|\mathbb{X}, \mathbb{I}) \cdot p(\mathbb{X}|\mathbb{I}) \cdot p(\mathbb{I}). \quad (1)$$

To build a multi-view pose estimation approach that can generalize to new environments, we make two conditionally independence assumptions. Firstly, the 3D pose Y is assumed conditionally independent of images \mathbb{I} given 2D poses \mathbb{X} . Obviously, this is not always correct, as one can find different 3D skeletons that have similar 2D projections due to the 3D-2D perspective effect. The likelihood of such cases, however, degrades dramatically in a multi-view setup, where a working volume has been captured from complementary views. Secondly, we assume that given an image observation for a view j , 2D poses in this view are conditionally independent of detections in the other views and other image observations. One can see that this assumption does not hold in case of occlusions. But, we believe that this assumption is reasonable for these three reasons: (1) there exist challenging single-view datasets, e.g. MS COCO and MPII Pose, which can be used to train robust single-view pose detection models; (2) recent deep neural network-based approaches have achieved very promising results on unseen data and reliably discriminate occluded joints from visible ones [10,34,35]; and (3) it yields an interesting modeling that allows us to train a 2D pose detector independently. Considering these two assumptions, we can rewrite the joint probability as:

$$p(Y, \mathbb{X}, \mathbb{I}) = p(Y|\mathbb{X}) \cdot \prod_{j=1}^V (p(X_j|I_j) \cdot p(I_j)). \quad (2)$$

This equation indicates that a 2D pose detector is applied in each view independently and that the 3D pose regression function is solely dependent upon 2D pose detections. We model the first term using a multi-view 3D regression function, described in Sect. 3.4. The input for this function is provided by concatenating 2D detections for each individual person across all views, which is presented in Sect. 3.2. The second term is the single-view pose detector explained next.

3.1 Single-view 2D pose detector

The relaxation assumption mentioned above allows us to use arbitrary complex models to detect and localize 2D body poses given single-view images. We therefore use the deep convolutional network of [10] as single-view pose detector. This approach is currently the state-of-the-art approach for multi-person 2D pose estimation. In addition to its reliable multi-person pose estimation performance, the approach runs in nearly real time. Given an image, the model generates a set of 2D poses, where each body pose is specified by a collection of 18 body parts. For each body part, the model provides its pixel coordinate and a detection confidence. The confidence values are in range $[0, 1]$, where zero indicates undetected body parts.



Fig. 2 Person matching using epipolar geometry. A set of points and their corresponding epipolar lines are shown for a pair of images captured from two different viewpoints at the same time step. (Best seen in color)

3.2 Concatenating detections across all views

Given the detected poses per view, we need to find correspondences across the views. As we assume that the camera system is fully calibrated (i.e. both camera intrinsic and extrinsic parameters are available), we use epipolar geometry to find correspondences [19]. Let us assume that for each pair of cameras (C, C') the camera parameters are given with respect to the first one:

$$C = K[I|0] \quad \text{and} \quad C' = K'[R|t], \quad (3)$$

where K and K' are camera intrinsic parameters and $[A|b]$ indicates extrinsic parameters. We can compute the fundamental matrix F by:

$$F = K'^{-T} R K^T [K R^T t]_{\times}, \quad (4)$$

where $[b]_{\times}$ is the skew matrix operator. The fundamental matrix encapsulates all cameras parameters and allows us to compute the corresponding epipolar line for a point in the other view, as illustrated in Fig. 2.

Here, we use the fundamental matrix to compute average distances between detected skeletons for all pairs of views. This distance is computed for each possible pair of detections from two distinct views as the average distance between a subset of body joints detected in both skeletons. We collect 2D skeletons for each person across two views by computing the average distances between detected skeletons in one view and the corresponding epipolar lines of skeletons from the other view and by then finding disjoint pairs of skeletons with the lowest average distance. We exclude pairs for which the average distance is bigger than 20 pixels.³ We then use the matched skeletons to establish multi-view correspondences per individual person. If the distances between all skeletons in two pairs are less than 20 pixels, we join the pairs. In cases where a skeleton is shared among two pairs and distance between other skeletons in the pairs are larger than the

³ The image size is 480×640 pixels.

threshold, we only accept the pair with the smaller distance and ignore the other pair. One should note that despite the availability of the correspondences, we cannot use triangulation because inaccurate detections lead to high error in 3D and, more importantly, joints might be detected in less than two views, especially in cluttered environments. We therefore use a regression function to compute the 3D positions of the body joints.

To prepare the input for the regression function, we concatenate skeletons across all views. If a person is not detected in a view, we fill the corresponding entry with zeros. Each body part is represented by three channels: two channels indicating pixel location and the third channel indicating the detection confidence.

3.3 Training data generation

As mentioned in the introduction, we generate training samples by projecting 3D skeletons into 2D. The model can therefore be trained on data generated from existing datasets or any set of valid 3D poses. The projected 2D skeletons are computed based on the camera setup used at test time. Since the single-view 2D pose detector used at test time can provide noisy detections, the model needs to be trained on similar noisy detection data to be able to generalize. We therefore evaluate our 2D pose detector on the Human3.6M dataset, which contains both images and ground truth 2D poses, to characterize its performance. We use these evaluation results to design a normally distributed noise model for each body joint. This noise is used to perturb training data. We then compute the confidence for the joint as:

$$\text{conf} = \max\left(1 - \frac{w}{\lambda \cdot \sigma}, 0\right), \quad (5)$$

where w is the amount of additive noise, which is sampled from a normal distribution with zero mean and standard deviation σ , and λ is a coefficient. We use this coefficient to set the confidence of a joint to zero, i.e. undetected, based on the relative amount of added noise with respect to the standard deviation. We use the evaluation results of [10] on Human3.6M, presented in Sect. 4.3, to set these parameters. As shown by the experiments, perturbing trained data and incorporating the confidence value are important for the method to generalize well to unseen data.

3.4 Multi-view 3D regression function

As mentioned earlier, the regression function relies solely on the detections provided by the single-view 2D pose detector. In contrast to [14,30,47], we do not need to model a complex function to directly map image pixel intensities into body part locations in 3D. Similar to [31], we model the 3D regression




Fig. 3 Regression network architecture. The network consists of several stages. Each stage includes four fully connected (FC) layers and intermediate supervision is provided by computing an L_2 loss at the last FC layer in each stage. The network takes as input a vector of 2D poses concatenated across all views for each individual person, as presented in Sect. 3.2

function using a simple multi-stage multilayer neural network.

The illustration of the network architecture is shown in Fig. 3. The network consists of several stages, where each stage is made of four fully connected (FC) layers. The first stage takes the multi-view 2D detections as input, described in Sect. 3.2. Every stage in this network is trained to regress for the desired output. This provides intermediate supervision at each stage and automatically alleviates the problem of vanishing gradient that happens when there are many intermediate layers between the network input and output layers [10]. We can therefore build deep neural networks by stacking several stages. The stage-wise supervision is provided by computing the L_2 loss between the output of the last layer in each stage and the desired output (y^*):

$$\mathcal{L}_s = \frac{1}{N} \sum_{n=1}^N \|y_n^s - y_n^*\|_2^2, \quad (6)$$

where \mathcal{L}_s is the average loss computed over all N training samples used in this iteration and y_n^s is the output of the last layer at stage s for sample n . The network is optimized by computing the overall network loss as a sum of the losses from all S stages that is defined as:

$$\mathcal{L} = \sum_{s=1}^S \mathcal{L}_s. \quad (7)$$

Since we need to retrain the model for new multi-view setups, we use batch normalization in order to reduce sensitivity to network initialization and learning rate [22]. We have also used dropout to avoid overfitting [44] and rectified linear units to achieve non-linearity [33].

4 Experiments

In this section, we present the evaluation on two multi-view datasets and compare with state-of-the-art results.

4.1 Implementation details

We implement our approach using TensorFlow [1]. In each stage of the network, the size of the first and last layers are set based on the input and output dimensions and the size of the intermediate layers are set to 1024. Our network is trained using the Adam optimizer. We set the starting learning rate to 0.001 and use exponential decay. The batch size is set to 512 and we train our network for 200 epochs. We observe that the performance of the network reaches a plateau when more than three stages are used. We therefore use three-stage networks throughout our experiments. A forward pass takes less than 1ms on a 1080Ti GPU. We can therefore say that the computation time of our multi-view regression model is almost negligible compared to the use of the 2D detector.

4.2 Datasets

Human3.6M. Human3.6M is currently the largest multi-view human pose estimation dataset. The dataset includes around 3.6 million images collected from 15 actions performed by seven professional actors in a laboratory environment [23]. The actions have been recorded by a four-view RGB camera system and camera parameters, including both intrinsic and extrinsic parameters, are available. Full-body 3D ground truth annotations are generated using a motion capture system. Following the standard evaluation protocol used in the literature, five subjects (S1, S5, S6, S7, S8) are used for training and two subjects (S9, S11) for testing [11,31,36]. Mean per joint position error (MPJPE) in millimeter is used as evaluation metric and test results are collected per action.

Multi-view OR. The multi-view OR (MVOR) dataset is, to the best of our knowledge, the first multi-view pose estimation dataset that is generated from recordings in an uncontrolled environment. All activities in an operating room have been recorded for four days using a three-view camera system [26]. We have selected every 1500 multi-view frames if there is at least one persons in one of the views. The dataset has been manually annotated to provide both 2D and 3D upper-body poses. The dataset includes around 700 multi-view frames and 1100 persons. The presence of multiple persons and clutter make this dataset much more challenging than Human3.6M as can be seen in Fig. 1. To report 2D body part localization on this dataset, we use the probability of correct keypoints (PCK) metric that is commonly used for evaluating multi-person pose estimation [10,26]. MPJPE is used to report 3D body part localization.

4.3 2D detection results

In this section, we evaluate the 2D detection model of [10] on both datasets to assess its performance on such unseen data.

Table 2 2D MPJPE results of the 2D detector [10] on the train set of Human3.6M. For each body part, the average MPJPE in pixels is computed across all actions in Human3.6M per camera

Camera ID	Hip	Knee	Foot	Shoulder	Elbow	Wrist	Avg
54138969	16	15	14	7	11	16	13
55011271	13	8	10	6	8	10	9
58860488	15	13	12	7	11	18	13
60457274	16	9	10	7	10	11	11
Avg	15	11	11	7	10	14	11

Table 3 PCK results on MVOR. Body part detection results are reported for both Deep3DPS [26] and Cao et al. [10] using the PCK metric. Note that Deep3DPS has been finetuned on another dataset captured in the same OR and relies on both color and depth images

	Head	Shoulder	Elbow	Wrist	Hip	Avg
Deep3DPS [26]	93.4	77.0	71.5	73.7	69.1	76.9
Cao et al. [10]	92.8	90.1	75.6	75.9	58.9	78.6

In addition, we use the results on Human3.6M to model the characteristics of the 2D detector, which are required by our data generation model presented in Sect. 3.3.

In Table 2, we present the results of the single-view 2D pose detector [10] on the Human3.6M train set. We should note that the detector has not seen any data from this dataset during training. We use MPJPE in pixel to compute body part localization errors. The results for each body parts are reported per camera. The results for head and neck localizations are not presented as the annotation for these body parts are different between Human3.6M and MS COCO [27] that is used to train the detector. Note that the detector is applied on the whole image, i.e. no bounding box is provided, in contrast to previous work that relies either on ground truth [23,31,32] or on person detectors [46] to obtain bounding boxes. In total, 3% of the joints are not detected and the detector achieves the average MPJPE of 11 pixels. It is worth mentioning that the detector performs similarly on the test set. Table 3 presents the results of the 2D detector on the MVOR dataset. The model attains an average PCK of 78.9% on this dataset. We have also reported the performance of Deep3DPS [26], which is the state-of-the-art model on this dataset. In contrast to [10], which is trained on the RGB images of MS COCO, the Deep3DPS model uses both color and depth images and has been trained on MPI Pose and then finetuned on a single-view OR dataset. The 2D pose detector of [10] outperforms Deep3DPS. These results show that the detector achieves fairly promising results on both datasets even without finetuning. Comparing the performance of the 2D detector on these two datasets also indicates that the MVOR dataset is much more complex, as the number of undetected joints is much higher (21% vs. 3%).

For generating the training data, the evaluation results on the train set of Human3.6M, which are reported in Table 2, are used to set the parameters of the noise model. The train set from Human3.6M is chosen to avoid any overlap between train and test sets. The coefficient λ in (5) is set to two. As a result, 5% of the joints will be labeled as undetected, which is on par with the percentage of undetected joints in Human3.6M.

4.4 3D localization results

Human3.6M. As Human3.6M is a fairly new dataset and state-of-the-art results are mainly reported using single-view models, we compare our approach with recent state-of-the-art single- and multi-view models for 3D pose estimation on Human3.6M. For the sake of comparison, we have therefore trained a variant of our proposed regression function that relies solely on single-view input. Table 4 reports evaluation results of our approach with different configurations. Models that are relying on single-view input are denoted by SV and multi-view ones by MV. These models are trained either on ground truth (GT) 2D poses, Noisy GT 2D poses as described in Sect. 3.3 or on 2D detections provided by either [34] or [10] for comparison. Even though Human3.6M is a single-person dataset, note that in [36,46] the input images are cropped using bounding boxes around the persons and that the 2D pose detector models of [34] and [50] used in [11] and [31] are applied on bounding boxes around the persons obtained from ground truth.

Our single-view 3D pose regression model trained on 2D detection provided by [34] achieves the average localization error of 67.2 mm. We should note that our results for this model improve slightly over the results reported by [31] on the same experimental setup (67.5), where the same 2D pose detector trained on MPII Pose is used without any finetuning on Human3.6M. [31] showed that the results can be improved by finetuning the model on Human3.6M (62.9 vs. 67.5), which is in line with the results reported in [11]. However, in order to easily generalize to new environments, we do not finetune 2D pose detectors as this would require annotated data. Except the model <SV, [34]>, which uses the same 2D pose detector during both training and testing for the sake of fair comparison with [31], all our models have used 2D detections provided by [10] during testing.⁴ We should note that even though our single-view 3D regression model trained on the 2D detections provided by [34] performs better than other variants of our single-view model, we decide to use the model of [10] instead, as it is not restricted to bounding boxes and allows us to detect and estimate 2D body poses in *multi-person scenarios*, e.g. the MVOR dataset.

⁴ Please note that at test time 2D poses are detected using [10] even in case of models trained on GT poses, which is different from [31].

Table 4 3D MPJPE in millimeter on Human3.6M. Quantitative results of our approach with different configurations are reported and compared with the state-of-the-art

Setting	Direc.	Discuss	Eat	Greet	Phone	Photo	Pose	Purch.	Sit	SitD	Smoke	Wait	Walk	WalkD	WalkT	Avg
Tekin et al. [46]	102.4	147.2	88.8	125.3	118.0	182.7	112.4	129.2	138.9	224.9	118.4	138.8	126.3	55.1	65.8	125.0
Chen and Ramanan [11]	89.9	97.6	89.9	107.9	107.3	139.2	93.6	136.0	133.1	240.1	106.6	106.2	87.0	114.0	90.5	114.1
Pavlakos et al. [36]	67.4	71.9	66.7	69.1	72.0	77.0	65.0	68.3	83.7	96.5	71.7	65.8	74.9	59.1	63.2	71.9
Martinez et al. [31]	51.8	56.2	58.1	59.0	69.5	78.4	55.2	58.1	74.0	94.6	62.3	59.1	65.1	49.5	52.4	62.9
<SV, Newell et al. [34]>	53.4	58.6	62.1	63.2	86.2	83.3	56	58.1	81.2	101.2	68.4	64.1	67.4	51	54.2	67.2
<SV, Cao et al. [10]>	69.5	75.5	67.6	76.8	84.6	94.9	69.8	68.4	92.2	113.7	77.1	75.1	77.2	59.0	64.2	77.7
<SV, GT>	94.2	113.7	96.9	106.5	119.8	127.6	86.5	149.9	145.6	222.3	113.5	111.2	120.9	92.8	92.4	119.6
<SV, Noisy GT>	69.7	78.8	69.8	77.5	84.4	97.6	64.9	86.5	103.3	125.8	81.8	80.4	83.3	59.9	62.6	81.8
Pavlakos et al. [37]	41.2	49.27	42.8	43.5	55.6	46.9	40.3	63.7	97.6	119.9	52.1	42.7	41.8	51.9	39.4	56.9
<MV, Cao et al. [10]>	39.4	46.9	41.0	42.7	53.6	54.8	41.4	50.0	59.9	78.8	49.8	46.2	51.1	40.5	41.0	49.1
<MV, GT>	92.1	105.8	110.1	94.0	128.2	117.0	77.0	152.2	152.0	227.5	122.9	104.3	125.1	88.7	80.9	118.5
<MV, Noisy GT>	47.1	60.5	48.7	53.5	63.5	71.1	48.7	57.8	72.2	81.7	59.0	55.9	60.6	43.4	44.3	57.9

SV indicates that our model is trained on single-view input and MV indicates multi-view input. Our models are trained on 2D poses obtained from: [34]'s model, [10]'s model, 2D projections of 3D ground truths (GT) and our data generation method described in Section 3.2 (Noisy GT). <M, P> denotes our model settings where M is either SV or MV and P indicates the type of poses that are used during training

The evaluation results show that our single-view model trained on ground truth 2D poses and the model of [11] perform similarly. This indicates that our regression function that is trained on perfect GT data will eventually work similarly to the lookup table used in [11]. One can therefore conclude that if perfect 2D detections are obtained, a 2D-to-3D regression function or a lookup table would work similarly. But, the 2D detections are not perfect in practice. Therefore, by incorporating detection noise during training as described in Sect. 3.3, we have constructed a model <SV, Noisy GT> that could cope better with noisy detection (81.8 vs. 119.6). We observe that if we train the model on 2D detections from the same 2D detector used during testing, i.e. [10], average MPJPE is improved by only four millimeters. These results indicate that our data generation model presented in Sect. 3.3 has properly incorporated the detector's characteristics and our approach generalizes well to test data.

We have also presented the evaluation results of our multi-view regression function in Table 4. Training the model <MV, [10]> on 2D pose detections by the same detector model as the one used at test time achieves the average MPJPE of 49 millimeters, which outperforms [37]. This is the lower limit for MPJPE on Human3.6M, which can be obtained by our MV regression model using this single-view pose detector. During our experiments, we observe that even though our multi-view regression models have generally converged to lower training losses compared to single-view ones, both single-view and multi-view models trained on ground truth poses achieve similar performance (119.6 vs. 118.5). We believe that as the multi-view model is only trained on perfect ground truth 2D poses, it always expects the exact projections of a 3D pose in all views. But, since the 2D pose detector provides noisy detections, this is not always possible at test time. The last row shows the results of our multi-view regression model trained using 2D poses generated from 3D ground truth by incorporating the 2D detector's characteristics. We should note that even without finetuning the detector on Human3.6M this model performs similarly to [37], which has been trained on Human3.6M. This model also reduces the error by more than 50% compared to the same model trained on ground truth data only. Furthermore, the model has also improved the localization results by $\sim 30\%$ compared to the single-view model <SV, Noisy GT> indicating that this model has properly incorporated 2D body part locations across all views to regress for their 3D positions. These results also confirm our hypothesis that incorporating the characteristics of the detector during training enables developing models that are robust to the inaccuracies and failures of the detector at test time.

Multi-view OR In order to assess the ability of our approach to generalize to new multi-view environments, we evaluate the performance of our approach on the multi-view OR dataset.

We use the 3D poses from Human3.6M, the camera calibration parameters of MVOR and the data generation model described in Sect. 3.2 to train a multi-view 3D regression model. The evaluation results of this model on MVOR are presented in Table 5. We use 3D MPJPE in centimeter as evaluation metric. Following the convention in MVDeep3DPS [26], MPJPE is computed for the same set of body parts and is reported per number of supporting views. Our model has achieved the average MPJPE of 17 cm on this dataset. The results show a significant improvement in the localization of the body parts as the number of supporting view increases. The average MPJPE is improved by 12 cm for persons who are detected in three views compared to those who are only detected in one view. This clearly indicates the benefit of observing an environment from multiple complementary views and the ability of our regression model to leverage such data for predicting 3D body poses even when some body parts are invisible.

Table 5 also compares the performance of our model with the MVDeep3DPS model [26]. We should note that MVDeep3DPS requires both color and depth images in contrast to our approach that relies solely on color images. Our approach, which only uses Human3.6M data, improves the results over MVDeep3DPS, even though MVDeep3DPS is trained on an annotated dataset recorded in the same OR as the one used to capture MVOR. This evaluation results demonstrate that our approach can exploit existing datasets to easily generalize to new multi-view setups without any need for new annotations.

4.5 Qualitative results

In Figs. 4 and 5, we show qualitative results on both Human3.6M and MVOR.⁵ Each row shows a multi-view frame. The predicted 3D poses are shown in the last column and the overlaid 2D poses are obtained by projecting the 3D poses into the views. Figure 4 demonstrates the high-quality of the predicted 3D body poses. For example, the frame presented in the last row shows that our approach can successfully incorporate evidence across all views to localize the occluded body parts.

We also show some frames from the multi-view OR dataset in Fig. 5. As can be seen in this figure, this dataset is much more complex due to the similar appearance of the objects as well as the people and the presence of many objects and multiple persons in the scenes. Our approach predicts fairly accurate 3D body poses and always correctly detects the left and right side labels even though it has not seen any

⁵ Please note that for generating the qualitative images, the predicted 3D poses are transferred to the room reference frame using an offset computed as the relative difference between the neck location in the ground truth and the neck location in the predicted skeleton.

Table 5 MPJPE in centimeter on MVOR. Quantitative results of our approach are compared with MVDeep3DPS [26]. We follow the same convention as in MVDeep3DPS and report the results per number of supporting viewpoints for the same set of body parts

Part name	One view		Two views		Three views	
	MVDeep3DPS	Ours	MVDeep3DPS	Ours	MVDeep3DPS	Ours
Shoulder	19	13	15	8	10	5
Hip	27	20	23	15	17	11
Elbow	27	25	23	19	16	12
Wrist	32	34	25	28	18	16
Average	26	23	22	18	15	11

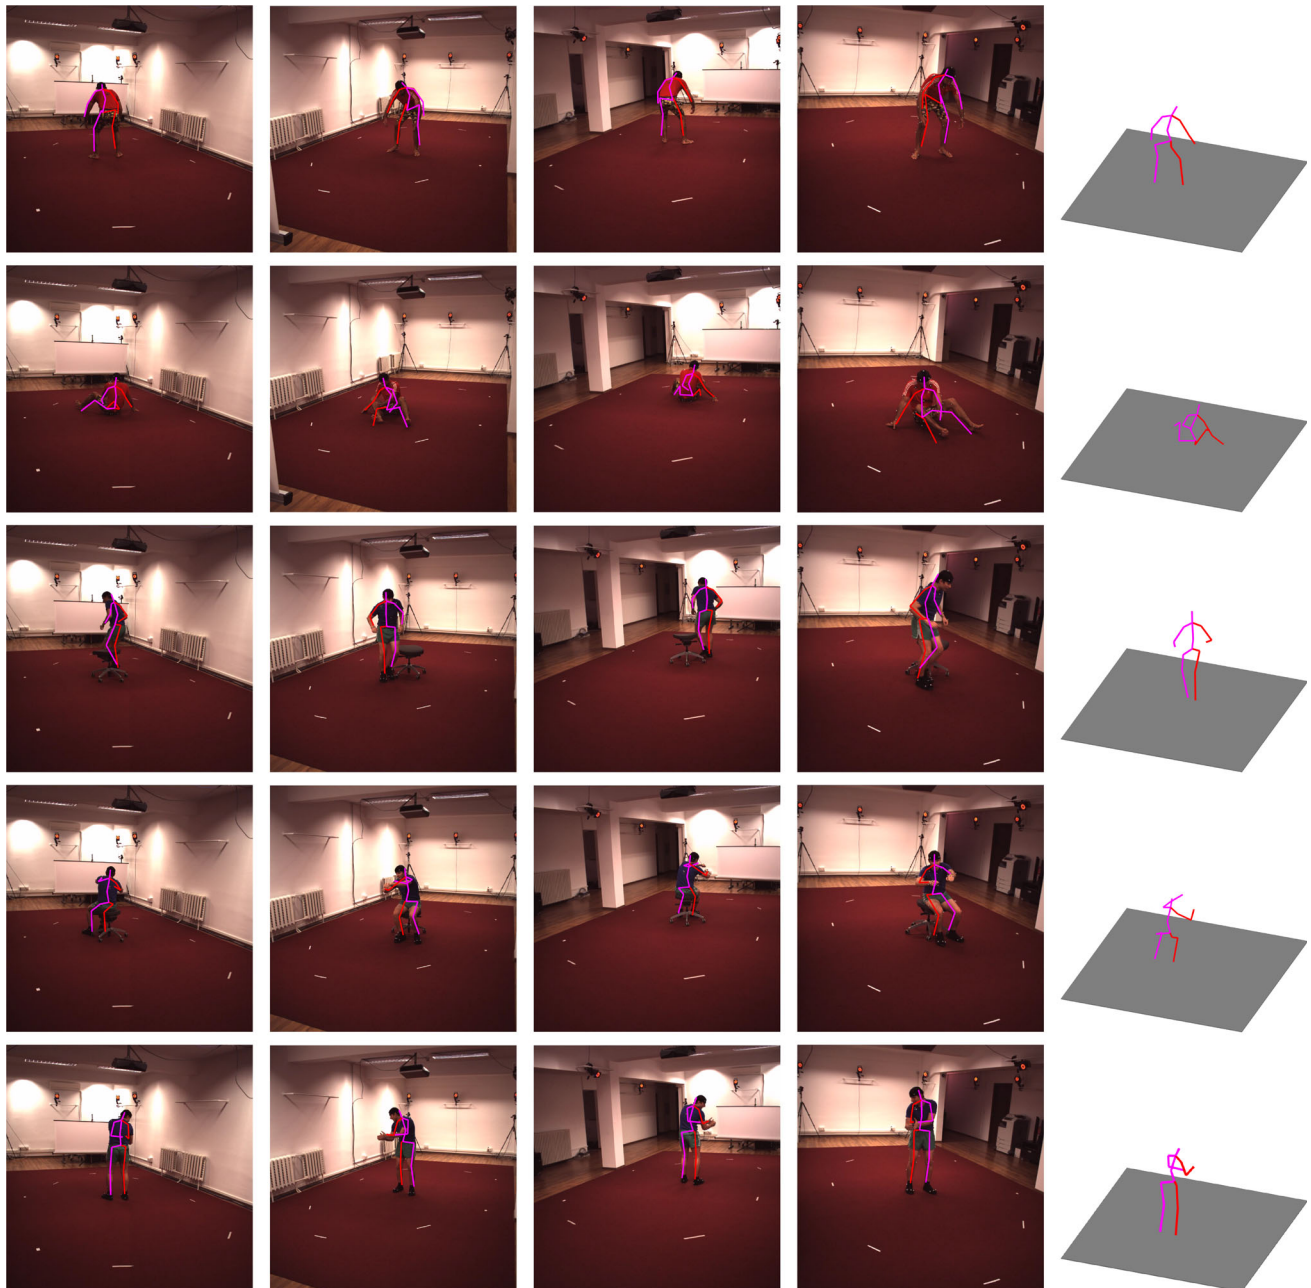


Fig. 4 Qualitative results on the Human3.6M dataset. The last column shows the 3D poses and the other columns show the corresponding multi-view frames where the projected 2D poses are superimposed. The body parts from the right side of the body are drawn in red. (Best seen in color)




Fig. 5 Qualitative results on the multi-view OR dataset. The first three columns in each row show a multi-view frame and the last column shows the corresponding 3D poses for that frame. The overlaid 2D skeletons

are computed by projecting 3D poses to the views. The body parts on the right side of the body are drawn in red. The blue arrow in the last row indicates a physically implausible body pose. (Best viewed in color)

data from this dataset or any other data collected in such an OR environment at the training stage.⁶

The complexity of this dataset also allows us to identify some of the limitations of the proposed approach. For example, we observe that the elbow and the wrist localization are less accurate compared to other body parts, which is in line with results presented in Tables 2 and 3. We envision that enforcing appearance consistencies among the projec-

tions of a body part across all views can be used to update and improve the 2D body joint detections. The improved 2D detections could then be fed into our multi-view regression model to obtain a more accurate localizations of the body parts in 3D. In the last row of Fig. 5, we have highlighted a 3D body pose, where the right arm configuration is infeasible because of body physical constraints. We believe that since our training data generation model described in Sect. 3.2 perturbs 3D poses randomly and does not take the body constraints into account, it may have generated such a train-

⁶ More qualitative results generated by our model on both datasets are available at https://youtu.be/Cx_kTRzqzA.

Table 6 Evaluation results on Human3.6M under noise. Ground truth data is used to both train and test the performance of a variant of our model that relies on single-view input. Similar to EDM [32] and SimpBase [31], we add noise to the test data and gradually increase the amount of noise, where $\mathcal{N}(0, \sigma)$ indicates a normal distribution with mean zero and standard deviation σ in pixel

	EDM	SimpBase	Ours
GT	62.2	37.1	47.2
GT+ $\mathcal{N}(0, 5)$	67.1	46.7	48.4
GT+ $\mathcal{N}(0, 10)$	79.1	52.8	50.8
GT+ $\mathcal{N}(0, 15)$	96.1	60.0	56.4
GT+ $\mathcal{N}(0, 20)$	115.6	70.2	65.7




Fig. 6 Data augmentation. The results of our multi-view model on MVOR are reported as a function of the number of random rotations of 3D body poses

ing sample. Therefore, it would be interesting to combine our data generation model with a model like the one used in [48] to enforce and verify the physical plausibility of the generated 3D poses.

4.6 Ablation study

We performed several experiments on Human3.6M to study the impact of each of the components of our approach. We first observe that by removing the stage-wise supervision, the performance always drops. For example, average MPJPE changes from 57.9 to 77.2 for our $\langle MV, \text{Noisy GT} \rangle$ model. Removing batch normalization leads to a substantial increase in the error (from 57.9 to 175). We also observe that the use of dropout during the training of single-view models and multi-view models on perfect ground truth data is important to obtain more robust models, as it reduces the errors by 20 – 50 mm. However, deactivating dropout for our multi-view models trained on [10]’s detections or Noisy GT decreases localization errors by 2 and 9 mm, respectively. We believe that this is due to the fact that 2D detection inputs are constructed from single-view poses that have been indepen-

dently affected by noise in each view by either the detector inaccuracy or by our data generation model. This independent noise can therefore work as a regularizer to enforce neurons to detect the most relevant information across all views, thereby removing the need for dropout.

Following [32] and [31], we perform a series of experiments to evaluate the performance of our approach under different levels of noise at test time. For a fair comparison, we evaluate our single-view model trained on Noisy GT and add different levels of Gaussian noise to ground truth 2D poses at test time. The evaluation results are presented in Table 6 and are compared with EDM [32] and SimpBase [31]. Even though the average localization error of SimpBase is lower than our model’s error by one centimeter when tested on perfect ground truth 2D poses, our model achieves lower localization errors as the noise increases. This indicates that incorporating the detector’s characteristics during training allows our model to better cope with the noise at test time.

In a multi-view setup, a 3D body pose can have completely different projections to the views depending on the orientation of the person with respect to the reference coordinate system. We therefore need to construct our multi-view regression model in a way that is robust to these changes in the orientation of the person, as our model only relies on these 2D projections to compute 3D body poses. For this reason, we propose to augment the training data by rotating each 3D pose in Human3.6M w.r.t. the reference frame. Figure 6 shows the effect of this data augmentation. We report the results of our multi-view model $\langle MV, \text{Noisy GT} \rangle$ on the MVOR dataset as a function of the number of rotations applied to each 3D poses in Human3.6M. The results show that applying up to three random rotations decrease the error but applying more random rotation does not lead to any improvement. Apart from the evaluation results reported in Fig. 6, for all the other evaluation on MVOR we always use our multi-view model trained on the train set of Human3.6M, which is augmented by applying three random rotations to each 3D pose.

5 Conclusions

We have presented an easily generalizable approach for estimating 3D body poses using multi-view data. We have proposed a two-step framework to tackle this problem, which separates single-view pose detection from multi-view 3D pose regression. The proposed approach permits to effectively exploit existing datasets to generalize to new multi-view environments. We have used a multi-stage neural network as regression function to estimate 3D poses. Our model has been trained on data generated from a set of valid 3D poses by projecting the 3D poses using the camera parameters used at the test time and by incorporating the characteristics of the single-view pose detector. Our evalua-

tion results have indicated the effectiveness and importance of incorporating the detector's characteristics during training, as it significantly reduced the localization error and achieved results on par with models trained on the output of the detector. We have also evaluated the generalization of our approach on the multi-person MVOR dataset by using only the camera configuration parameters from this dataset during training, but no image data. Our approach yielded fairly accurate results and outperforms the state-of-the-art model on this dataset. The results also showed that the localization error dramatically decreases as the number of supporting views increases. This highlights the benefit of our approach in leveraging multi-view data to obtain a reliable model for crowded and cluttered environments. To the best of our knowledge, this is also the first multi-view RGB approach that has been quantitatively evaluated on a real-world dataset for the task of 3D body part localization.

Future work will focus on improving the training data generation approach by taking human body kinematic constraints into account and on extending the 3D regressor model to enforce appearance similarity among different views.

Acknowledgements This work was supported by French state funds managed within the Investissements d'Avenir program by BPI France (project CONDOR) and by the ANR (references ANR-11-LABX-0004, ANR-10-IAHU-02 and ANR-16-CE33-0009). The authors would also like to acknowledge the support of NVIDIA with the donation of a GPU used in this research.

References

- Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G.S., Davis, A., Dean, J., Devin, M., Ghemawat, S., Goodfellow, I., Harp, A., Irving, G., Isard, M., Jia Y., Jozefowicz, R., Kaiser, L., Kudlur, M., Levenberg, J., Mané, D., Monga, R., Moore, S., Murray, D., Olah, C., Schuster, M., Shlens, J., Steiner, B., Sutskever, I., Talwar, K., Tucker, P., Vanhoucke, V., Vasudevan, V., Viégas, F., Vinyals, O., Warden, P., Wattenberg, M., Wicke, M., Yu, Y., Zheng, X.: TensorFlow: large-scale machine learning on heterogeneous systems, 2015. URL <https://www.tensorflow.org/>. Software available from tensorflow.org
- Amin, S., Andriluka, M., Rohrbach, M., Schiele, B.: Multi-view pictorial structures for 3d human pose estimation. In: British Machine Vision Conference (BMVC), September (2013)
- Amin, S., Müller, P., Bulling, A., Andriluka, M.: Test-time adaptation for 3d human pose estimation. *Pattern Recogn.* **8753**, 253–264 (2014)
- Andriluka, M., Pishchulin, L., Gehler, P., Schiele, B.: 2D human pose estimation: New benchmark and state of the art analysis. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 3686–3693, (2014)
- Belagiannis, V., Amin, S., Andriluka, M., Schiele, B., Navab, N., Ilic, S.: 3d pictorial structures for multiple human pose estimation. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 1669–1676, (2014)
- Belagiannis, V., Wang, X., Schiele, B., Fua, P., Ilic, S., Navab, N.: Multiple human pose estimation with temporally consistent 3D pictorial structures. In: ChaLearn Looking at People Workshop, European Conference on Computer Vision (ECCV2014), pp. 742–754, September (2014)
- Belagiannis, V., Wang, X., Shitrit, H.B.B., Hashimoto, K., Stauder, R., Aoki, Y., Kranzfelder, M., Schneider, A., Fua, P., Ilic, S., Feussner, H., Navab, N.: Parsing human skeletons in an operating room. *Machine Vision and Applications*, pp. 1–12, (2016)
- Biswas, P., Liang, T.C., Toh, K.C., Ye, Y., Wang, T.C.: Semidefinite programming approaches for sensor network localization with noisy distance measurements. *IEEE Trans. Autom. Sci. Eng.* **3**(4), 360–371 (2006)
- Burenius, M., Sullivan, J., Carlsson, S.: 3d pictorial structures for multiple view articulated pose estimation. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 3618–3625, (2013)
- Cao, Z., Simon, T., Wei, S.-E., Sheikh, Y.: Realtime multi-person 2D pose estimation using part affinity fields. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 1302–1310, (2017)
- Chen, C.-H., Ramanan, D.: 3D human pose estimation = 2D pose estimation + matching. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 5759–5767, July (2017)
- Chen, D., Chou, P.-C., Fookes, C.B., Sridharan, S.: Multi-view human pose estimation using modified five-point skeleton model. In: International Conference on Signal Processing and Communication Systems, pp. 17–19 (2008)
- Dogan, E., Eren, G., Wolf, C., Lombardi, E., Baskurt, A.: Multi-view pose estimation with mixtures-of-parts and adaptive view-point selection. In: IET Computer Vision (2017)
- Fang, H., Xu, Y., Wang, W., Liu, X., Zhu, S.C.: Learning knowledge-guided pose grammar machine for 3d human pose estimation. *CoRR*, abs/1710.06513, 2017. URL <http://arxiv.org/abs/1710.06513>
- Felzenszwalb, P.F., Huttenlocher, D.P.: Pictorial structures for object recognition. *International Journal of Computer Vision* **61**(1), 55–79 (2005)
- Felzenszwalb, P.F., Girshick, R.B., McAllester, D., Ramanan, D.: Object detection with discriminatively trained part-based models. *IEEE Trans. Pattern Anal. Mach. Intell.* **32**(9), 1627–1645 (2010)
- Fischler, M.A., Elschlager, R.A.: The representation and matching of pictorial structures. *IEEE Trans. Comput.* **22**(1), 67–92 (1973)
- Gall, J., Rosenhahn, B., Brox, T., Seidel, H.-P.: Optimization and filtering for human motion capture. *Int. J. Comput. Vis.* **87**(1), 75–92 (2010)
- Hartley, R.I., Zisserman, A.: *Multiple View Geometry in Computer Vision*. Cambridge University Press, Cambridge (2000)
- Hofmann, M., Gavrilu, D.M.: Multi-view 3D human pose estimation in complex environment. *Int. J. Comput. Vis.* **96**(1), 103–124 (2011)
- Huang, F., Zeng, A., Liu, M., Lai, Q., Xu, Q.: Deepfuse: an IMU-aware network for real-time 3d human pose estimation from multi-view image. In: The IEEE Winter Conference on Applications of Computer Vision (WACV) (2020)
- Ioffe, S., Szegedy, C.: Batch normalization: accelerating deep network training by reducing internal covariate shift. In: International Conference on Machine Learning, pp. 448–456. (2015)
- Ionescu, C., Papava, D., Olaru, V., Sminchisescu, C.: Human3.6m: Large scale datasets and predictive methods for 3d human sensing in natural environments. *IEEE Trans. Pattern Anal. Mach. Intell.* **36**(7), 1325–1339 (2014)
- Jiang, H.: 3d human pose reconstruction using millions of exemplars. In: International Conference on Pattern Recognition, pp. 1674–1677, (Aug 2010)
- Jiang, M., Zhuliang, Y., Zhang, Y., Wang, Q., Li, C., Lei, Y.: Reweighted sparse representation with residual compensation for 3d human pose estimation from a single rgb image. *Neurocomputing* **358**, 332–343 (2019)

26. Kadkhodamohammadi, A., Gangi, A., de Mathelin, M., Padoy, N.: A multi-view RGB-D approach for human pose estimation in operating rooms. In: Proceedings of IEEE Winter Conference on Applications of Computer Vision (WACV), pp. 363–372, (2017)
27. Lin, T.-Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft COCO: CommonLawrence: Microsoft COCO: Context, pp. 740–755. Springer (2014)
28. Luo, X., Berendsen, B., Tan, R.T., Veltkamp, R.C.: Human pose estimation for multiple persons based on volume reconstruction. In: International Conference on Pattern Recognition, pp. 3591–3594 (2010)
29. Luvizon, D., Picard, D., Tabia, H.: Multi-task deep learning for real-time 3d human pose estimation and action recognition. In: IEEE Transactions on Pattern Analysis and Machine Intelligence (2020)
30. Luvizon, Diogo C., Picard, David, Tabia, Hedi: 2D/3D pose estimation and action recognition using multitask deep learning. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR), (2018)
31. Martinez, J., Hossain, R., Romero, J., Little, J.J.: A simple yet effective baseline for 3d human pose estimation. In: IEEE International Conference on Computer Vision (ICCV), pp. 2659–2668, (2017)
32. Moreno-N.: Francesc: 3d human pose estimation from a single image via distance matrix regression. In: 2017 IEEE Conference on Computer Vision and Pattern Recognition, (CVPR), pp. 1561–1570, (2017)
33. Nair, V., Hinton, G.E.: Rectified linear units improve restricted boltzmann machines. In: International Conference on Machine Learning, pp. 807–814, (2010)
34. Newell, A., Yang, K., Deng, J.: Stacked Hourglass Networks for Human Pose Estimation, pp. 483–499. (2016)
35. Newell, A., Huang, Z., Deng, J.: Associative embedding: End-to-end learning for joint detection and grouping. *Advances in Neural Information Processing Systems* **30**, 2277–2287 (2017)
36. Pavlakos, G., Zhou, X., Derpanis, K.G., Daniilidis, K.: Coarse-to-fine volumetric prediction for single-image 3D human pose. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 1263–1272, (2017)
37. Pavlakos, G., Zhou, X., Derpanis, K.G., Daniilidis, K.: Harvesting multiple views for marker-less 3d human pose annotations. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 1253–1262, (2017)
38. Radwan, I., Dhall, A., Goecke, R.: Monocular image 3D human pose estimation under self-occlusion. In: International Conference on Computer Vision (ICCV), pp. 1888–1895, (2013)
39. Rogez, G., Weinzaepfel, P., Schmid, C.: LCR-Net: Localization-classification-regression for human pose. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 1216–1224, (2017)
40. Sigal, L., Balan, A.O., Black, M.J.: Humaneva: Synchronized video and motion capture dataset and baseline algorithm for evaluation of articulated human motion. *International Journal of Computer Vision* **87**(1), 4–27 (2009)
41. Sigal, L., Isard, M., Haussecker, H., Black, M.J.: Loose-limbed people: Estimating 3D human pose and motion using non-parametric belief propagation. *Int. J. Comput. Vis.* **98**(1), 15–48 (2012)
42. Slembrouck, M., Luong, H., Gerlo, J., Schütte, K., Van Cauwe-laert, D., De Clercq, D., Vanwanseele, B., Veelaert, P., Philips, W.: Multiview 3d markerless human pose estimation from openpose skeletons. In: Jacques B.-T., Patrice D., Wilfried P., Dan P., Paul S. (eds), *Advanced Concepts for Intelligent Vision Systems*, pp. 166–178. Springer (2020)
43. Srivastav, V., Issenhuht, T., Kadkhodamohammadi, A., de Mathelin, M., Gangi, A., Padoy, N.: MVOR: a multi-view RGB-D operating room dataset for 2D and 3D human pose estimation. *CoRR* (2018). URL <http://arxiv.org/abs/1808.08180>
44. Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., Salakhutdinov, R.: Dropout: a simple way to prevent neural networks from overfitting. *J. Mach. Learn. Res.*, 15(1):1929–1958 (2014). ISSN 1532-4435
45. Sundaresan, A., Chellappa, R.: Multicamera tracking of articulated human motion using shape and motion cues. *IEEE Trans. Image Process.* **18**(9), 2114–2126 (2009)
46. Tekin, B., Rozantsev, A., Lepetit, V., Fua, P.: Direct prediction of 3d body poses from motion compensated sequences. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 991–1000 (2016)
47. Tekin, B., Katircioglu, I., Salzmann, M., Lepetit, V., Fua, P.: Structured prediction of 3d human pose with deep neural networks. In: Proceedings of the British Machine Vision Conference BMVC (2016)
48. Vondrak, M., Sigal, L., Jenkins, O.C.: Physical simulation for probabilistic motion tracking. In: 2008 IEEE Conference on Computer Vision and Pattern Recognition, pp. 1–8 (2008)
49. Wan, Q., Zhang, W., Xue, X.: Deepskeleton: skeleton map for 3D human pose regression. *CoRR*, [arXiv:1711.10796](https://arxiv.org/abs/1711.10796) (2017). URL <http://arxiv.org/abs/1711.10796>
50. Wei, S.-E., Ramakrishna, V., Kanade, T., Sheikh, Y.: Convolutional pose machines. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 4724–4732 (2016)
51. Yu, J., Hong, C.: Exemplar-based 3d human pose estimation with sparse spectral embedding. *Neurocomputing* **269**, 82–89 (2017)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.