This ACCV 2020 paper, provided here by the Computer Vision Foundation, is the author-created version. The content of this paper is identical to the content of the officially published ACCV 2020 LNCS version of the paper as available on SpringerLink: https://link.springer.com/conference/accv



A Global to Local Double Embedding Method for Multi-person Pose Estimation

Yiming $Xu^{1[0000-0002-2013-3973]}$, Jiaxin $Li^{2[0000-0002-4790-3173]}$, Yan $Ding^{2*[0000-0002-9620-8679]}$, and Hua-Liang Wei^{3[0000-0002-4704-7346]}

¹ Yingcai Honors College, University of Electronic Science and Technology of China, Chengdu, China

² Key Laboratory of Dynamics and Control of Flight Vehicle, Ministry of Education, School of Aerospace Engineering, Beijing Institute of Technology, Beijing 100081,

China

³ Department of Automatic Control and Systems Engineering, University of Sheffield, Sheffield S1 3JD, UK

Abstract. Multi-person pose estimation is a fundamental and challenging problem to many computer vision tasks. Most existing methods can be broadly categorized into two classes: top-down and bottom-up methods. Both of the two types of methods involve two stages, namely, person detection and joints detection. Conventionally, the two stages are implemented separately without considering their interactions between them, and this may inevitably cause some issue intrinsically. In this paper, we present a novel method to simplify the pipeline by implementing person detection and joints detection simultaneously. We propose a Double Embedding (DE) method to complete the multi-person pose estimation task in a global-to-local way. DE consists of Global Embedding (GE) and Local Embedding (LE). GE encodes different person instances and processes information covering the whole image and LE encodes the local limbs information. GE functions for the person detection in top-down strategy while LE connects the rest joints sequentially which functions for joint grouping and information processing in A bottom-up strategy. Based on LE, we design the Mutual Refine Machine (MRM) to reduce the prediction difficulty in complex scenarios. MRM can effectively realize the information communicating between keypoints and further improve the accuracy. We achieve the competitive results on benchmarks MSCOCO, MPII and CrowdPose, demonstrating the effectiveness and generalization ability of our method.

1 Introduction

Human pose estimation aims to localize the human facial and body keypoints (e.g., nose, shoulder, knee, etc.) in the image. It is a fundamental technique for many computer vision tasks such as action recognition [1], human-computer interaction [2, 3], person Re-ID [4] and so on.

^{*}First two authors equally contributed to this work. Yan Ding is the corresponding author: dingyan@bit.edu.cn

Most of the existing methods can be broadly categorized into two classes: top-down methods [5–13] and bottom-up methods [14–18]. As shown in Figure 1 (a), the top-down strategy first employs a human detector to generate person bounding boxes, and then performs single person pose estimation on each individual person. On the contrary, the bottom-up strategy locates all body joints in the image and then groups joints to corresponding persons. Top-down strategy is less efficient because the need to perform single person pose estimation on each detected instance sequentially. Also, the performance of top-down strategy is highly dependent on the quality of person detections. Compared to top-down strategy, the complexity of bottom-up strategy is independent of the number of people in the image, which makes it more efficient. Though as a faster and more likely to be the real-time technique, the bottom-up methods may suffer from solving an NP-hard graph partition problem [15, 14] to group joints to corresponding persons on densely connected graphs covering the whole image.



Fig. 1. Comparison between (a) existing two-stage strategy and (b) our Double Embedding method for multi-person pose estimation. The proposed DE model implement the person detection and joints detection parallel, overcoming the intrinsic problems of existing two-stage based top-down and bottom-up strategies.

We analyze and try to bypass the disadvantages of these two conventional strategies. The low efficiency of the top-down strategy comes from the independent single person pose estimation on each person bounding box. For bottom-up strategy, treating all the detected unidentified joints equally causes the high difficulty of joints grouping process. Both top-down and bottom-up strategy are

3

two-stage structure with little interaction between the two stages. They both suffer from the separation of person instance detection and joints detection.

To overcome this intrinsic limitation, we propose to implement person detection and joints detection simultaneously, and realize information communicating between the two procedures to better utilize the structure information of human body. The proposed approach is illustrated in Figure 1 (b). We observe that the torso joints (shoulders and hips) are more stable than other limbs joints. With much lower degree of freedom than limbs joints, torso joints can represent the identity information to distinguish the different human instances well. We also introduce the center joint of the human body. Center joint is calculated by the average location of annotated joints. Together with the center joint point, torso joints and center points compose the Root Joints Group (RJG). Based on this, we categorize the rest joints on limbs into Adjacency Joints Group (AJG). In this paper, we propose the Double Embedding method to simplify the pipeline and improve the joints localization accuracy. The framework of the DE approach is shown in Figure 2. Double Embedding consists of Global Embedding (GE) and Local Embedding (LE). GE functions for the person instance separation process through encoding the clustering information of RJG. We follow the associative embedding method [18] to allocate 1D tags to each pixel related to the joints in root joints group. Joints belong to the same person have similar tags while joints belong to different instances have dissimilar tags.

GE encodes global information around the whole image, while LE focuses on local information of each instance based on the global clues from GE. AGJ is connected to identified RJG by corresponding displacement vector field encoded in LE. Basically, we take the center joint as the reference point. However, the displacements from extremities joints (wrists and ankles) to center joint are long-range displacements which are vulnerable to background noises. To optimize the long-range displacement prediction [19], we further divide AGJ into two hierarchies: the first level consists of elbows and knees, the second level consists of wrists and ankles. AJG is connected sequentially from the second level to first level and finally to torso joints in RJG which are identified. Long-range displacements are factorized into accumulative short-range displacements targeting on torso joints (hip joints and shoulder joints). Take the left ankle for example, the displacement from ankle to center joint is long-range displacement which is difficult to predict. To tackle this problem with better utilizing the body structure information, we change the reference joint to left hip, and divide the displacement from left ankle to left hip into shorter displacements: the displacement from left ankle to left knee and the displacement from left knee joint to left hip joint. Thus, AJG (limbs joints) are connected to RJG and identified in sequence. As for facial joints (e.g., Eves, nose, etc.), we localize them from predicted heatmaps and connect them with the long-range displacements targeting on the center joint.

In addition, we design Mutual Refine Machine (MRM) to further improve the joints localization accuracy and reduce the prediction difficulty in complex scenarios such as pose deformation, cluttered background, occlusion, person over-

lapping and scale variations. Based on hierarchical displacements and connection information encoded in LE, MRP refines the poor-quality predicted joints by high-quality predicted neighboring joints.

We reduce the person detection task to identifying RJG with associative embedding. This is essential to implement the person detection and joints detection at the same time. This is essential to implement the person detection and the following joints detection and grouping at the same time. Avoiding the independent single person pose estimation on each detected person bounding boxes makes the method more efficient. Compared to directly processing all the unidentified joints around the whole image, LE performs local inference with robust global affinity cues encoded in GE, reducing complexity for joints identifying. Different with the independence of two stages in previous two-stage strategy, GE and LE works mutually to complete the person detection and joints detection parallel.

We implement DE with Convolutional Neural Networks (CNNs) based on the state-of-the-art HigherHRNet[20] architecture. Experiments on benchmarks MSCOCO [21], MPII [22] and CrowdPose [23] demonstrate the effectiveness of our method.

The main contributions of the paper is summarized as follows:

- We attempt to simplify the pipeline for multi-person pose estimation, solving the task in global-to-local way.
- We propose the Double Embedding method to implement person detection and joints detection parallel, overcoming the intrinsic disadvantages caused by two-stage structure.
- Our model achieves competitive performance on multiple benchmarks.

2 Related Works

Top-Down Multi-Person Estimation Top-down methods [5–13] first employ object detection [24–27] to generate person instances within person bounding boxes, and then detect the keypoints of each person by single person pose estimation independently.

Mask R-CNN [9] adopts a branch for keypoints detection based on Faster R-CNN [24]. G-RMI [8] directly divides top-down methods as two stages and employs independent models for person detection and pose estimation. In [28], Gkioxari et al. adopted the Generalized Hough Transform framework for person instance detection, and then classify joint candidates based on the poselets. In [29], Sun et al. proposed a part-based model to jointly detect person instances and generate pose estimation. Recently, both person detection and single person pose estimation benefit a lot from the thriving of deep learning techniques. Iqbal and Gall [16] adopted Faster-RCNN [24] for person detection and convolutional pose machine [30] for joints detection. In [11], Fang et al. used spatial transformer network [31] and Hourglass network [13] for joints detection.

Though these methods have achieved excellent performance, they suffer from high time complexity due to sequential single person pose estimation on each

5

person proposal. Differently, DE performs the person detection and joints detection parallel, which simplifies the pipeline.

Bottom-Up Multi-Person Pose Estimation Bottom-up methods [14–18] detect all the unidentified joints in an image and then group them to corresponding person instances.

Openpose [32] proposes part affinity field to represent the limbs. The method calculates line integral through limbs and connects joints with the largest integral. In [18], Newell et al. proposed associate embedding to assign each joint with a 1D tag and then group joints which have the similar tags. PersonLab [19] groups joints by a 2D vector field in the whole image. In [33], Sven Kreiss et al. proposed Part Intensity Field (PIF) to localize body parts and Part Association Field (PAF) to connect body parts to form full human poses.

Nevertheless, the joints grouping cues of all these methods are covering the whole image, which causes high complexity for joints grouping. Different with the prior methods, global clues from GE reduce the search space for graph partition problem, avoiding high complexity of joint partition in bottom-up strategy.



Fig. 2. Overview of the Double Embedding (DE) model. For an image, we generate two kinds of feature maps for Global Embedding (GE) and Local Embedding (LE). GE and LE works parallel with information communicating to support each other. Based on GE and LE, we design Mutual Refine Machine (MRM) to refine the low-quality predicted joints which further improves the accuracy.

3 Double Embedding Method

In this section, we present our proposed Double Embedding method. Figure 2 illustrates the overall architecture of the proposed approach.

Joints Feature Map For an image I, we generate two kinds of feature maps from backbone network, one for Global Embedding (GE) and one for Local Embedding (LE). We use $J_R = \{J_{R1}, J_{R2}, ..., J_{RU}\}$ to denote Root Joints Group (RJG), J_{Ri} is the i-th kind of joints in RJG for all N persons in image I, and U is the number of joint categories in RJG. Similarly, we use $J_A = \{J_{A1}, J_{A2}, ..., J_{AV}\}$ to denote Adjacency Joints Group (AJG), J_{Ai} is the i-th kind of joints in AJG for all persons in image I, and V is the number of joint categories in RJG. For Global Embedding, let $h_{Gk} \in R^{W \times H}(k = 1, 2, ..., U)$ denote the feature map for the k-th kind of joint in Root Joints Group. For Local Embedding, $h_{Lk} \in R^{W \times H}(k = 1, 2, ..., V)$ denotes the feature map for the k-th kind of joint in AJG. The form and generation method of h_{Gk} and h_{Lk} are the same. To simplify the description, we use h_{fk} to represent both h_{Gk} and h_{Lk} .

It was pointed out that Directly regressing the absolute joint coordinates in an image is difficult[34, 14]. We therefore use heatmap, a confidence map models the joints position as Gaussian peaks. For a position (x, y) in image I, $h_{fk}(x, y)$ is calculated by:

$$h_{fk}(x,y) = \begin{cases} exp(-\frac{\|(x,y) - (x_k^i, y_k^i)\|_2^2}{\sigma^2}), & (x,y) \in \aleph_k^i \\ 0, & \text{otherwhise} \end{cases}$$
(1)

In which σ is an empirical constant to control the variance of Gaussian distribution, set as 7 in our experiments. (x_k^i, y_k^i) denotes the i-th groundtruth joint position in h_{fk} . $\aleph_k^i = (x, y) | ||(x, y) - (x_k^i, y_k^i)||_2 \leq \tau$ is regressing area for each joint to truncate the Gaussian distribution. Thus, we generate two kinds of feature map for GE and LE. Joints location of RJG and AJG are derived through NMS process.

Global Embedding Global Embedding functions as a simpler person detection process, reduces person detection problem to identifying the RJG. We use Associate Embedding for this process. The identification information in h_{Gk} is encoded in 1D tag space T_k for k-th joint in RJG. Based on the pixel location derived from peak detections in feature map h_{Gk} , corresponding tags are retrieved at the same pixel location in T_k . Joints belonging to one person have similar tags while tags of joints in different persons have obvious difference. Let $p = \{p_1, p_2, ..., p_N\}$ denote the N persons containing in image **I**. GE can be represented as a function $f_{GE} : h_{Gk} \to T_k$. f_{GE} learns to densely transform every pixel in h_{Gk} to embedding space T_k . We use $loc_{(p_n, J_{Rk})}(n = 1, 2, ..., N, k = 1, 2, ..., U)$ to denote the ground truth pixel location of the k-th kind of joint in RJG of n-th person.

If U' joints are labeled, the reference embedding for the n-th person is the average of retrieved tags of RJG in this person:

$$\bar{Tag}_n = \frac{1}{U'} \sum_k T_k(loc_{(p_n, J_R k)})$$
⁽²⁾

To pull the tags of joints within an individual together, pull-loss computes the squared distance between the reference embedding and the predicted embedding for each joint.

$$L_{pull} = \frac{1}{NU'} \sum_{n} \sum_{k} (\bar{Tag}_n - T_k(loc_{(p_n, J_{Rk})}))^2$$
(3)

To push the tags of joints in different persons, push-loss penalizes the reference embeddings that are close to each other. As the distance between two tags increases, push-loss drops exponentially to zero resembling probability density function of Gaussian distribution:

$$L_{push} = \frac{1}{N^2} \sum_{n} \sum_{n'} exp\{-\frac{1}{2\sigma^2} (\bar{Tag}_n - \bar{Tag}_{n'})^2\}$$
(4)

The loss to train the model as f_{GE} is the sum of L_{pull} and L_{push} :

$$L^{G} = \frac{1}{NU'} \sum_{n} \sum_{k} (\bar{Tag}_{n} - T_{k}(loc_{(p_{n},J_{Rk})}))^{2} + \frac{1}{N^{2}} \sum_{n} \sum_{n'} exp\{-\frac{1}{2\sigma^{2}}(\bar{Tag}_{n} - \bar{Tag}_{n'})^{2}\}$$
(5)

Local Embedding Local Embedding performs local inference and builds the connection clues between AJG and identified RJG. Relative position information of h_{Lk} is encoded in displacement space D_k for k-th joint in AJG. For each joint in AJG, we get its corresponding normed displacement to RJG from D_k at the same position in h_{Lk} . We first build the basic displacement to connect to the center joint. Basic displacement for the k-th joint in AJG of the n-th person is represented as the 2D vector:

$$Dis_{n}^{k} = (x_{n}^{r}, y_{n}^{r}) - (x_{n}^{k}, y_{n}^{k})$$
(6)

In which, (x_n^r, y_n^r) is the location of the center joint in the n-th person. Besides, we design the hierarchical displacement to connect the limbs joints to corresponding torso joints. Compared to basic displacements directly targeting on the center joint, hierarchical displacements are shorter ones which are more robust and easier to predict. Normally, we use hierarchical displacements in inference process. But if some intermediate joints are absent, we directly use the long-range prediction to complete the inference.

The general displacement for joint A to joint B of n-th person is:

$$Dis_{n}^{A2B} = (x_{n}^{B}, y_{n}^{B}) - (x_{n}^{A}, y_{n}^{A})$$
(7)

In some cases, we may use the property $Dis_n^{B2A} = -Dis_n^{A2B}$ to get reverse displacements of paired joints.

Local Embedding $f_{LE} : h_{Lk} \to D_k$ maps each pixel in feature map h_{Lk} to the embedding space D_k . For learning f_{LE} , we build the target regression map T_n^A for the displacement vector from joint A (the k-th kind of joint in AJG) to joint B of the n-th person as follows:

$$D_k^A(x,y) = \begin{cases} Dis_n^{(x,y)2(x_n^B, y_n^B)}/Z, & \text{if}(x,y) \in \aleph_k^A\\ 0, & \text{otherwise} \end{cases}$$
(8)

where $\aleph_k^A = \{(x,y) | \| (x,y) - (x_n^A, y_n^A) \|_2 \leq \tau \}$. The displacements are created in \aleph_k^A which is the same as regression area in h_{Lk} . $Z = \sqrt{H^2 + W^2}$ is the normalization factor, with H and W denoting the height and width of the image.

The starting point A (x_n^A, y_n^A) is generated from peak detections in feature map h_{Lk} , we get its corresponding displacement to joint B in D_k as $Dis_n^{A2B} = D_k((x_n^A, y_n^A))$. The ending joint is obtained: $(x_n^{End}, y_n^{End}) = (x_n^A, y_n^A) + Z \cdot Dis_n^{A2B}$. Compared with the peak detections in h_{LB} (containing the same category joints as joint B for all persons in the image, including (x_n^B, y_n^B)), it will be confirmed that joint B is the ending joint of joint A. Accordingly, joint A is connected to joints B, meaning they share the same identification. In this way, joints in AJG are connected to RJG and identified.

Mutual Refine Machine We design the Mutual Refine Machine (MRM) to reduce the prediction difficulty in complex scenes. For a low-quality predicted joint, it can be refined by the neighboring high-quality joints. Based on the displacements and connection information in LE, MRM realizes the information communicating between paired joints. For n-th person, if prediction probability of i-th joint in k-th kind of joints confidence map $h_{fk}((x_n^i, y_n^i))$ is lower than its neighboring paired joints $\{h_{fk'}((x_n^{i'}, y_n^{i'}))\}$ (in which $h_{fk'}((x_n^{i'}, y_n^{i'})) > 0.75)$, then we refine the location of i-th joint with the weighted fusion of its neighboring joints. Refined location is:

$$(x_n^i, y_n^i)_{refined} = \frac{h_{fk}((x_n^i, y_n^i))}{Q} * (x_n^i, y_n^i) + \sum_{i'} \frac{h_{fk'}((x_n^{i'}, y_n^{i'}))}{Q} * ((x_n^{i'}, y_n^{i'}) + Dis_n^{i'2i})$$
(9)

$$Q = h_{fk}((x_n^{i'}, y_n^{i'})) + \sum_{i'} h_{fk'}((x_n^{i'}, y_n^{i'}))$$
(10)

Training and inference To train our model, we adopt L2 loss L^H for joint confidence regression, smooth L1 loss [35] L^D for displacements regression and L^G for GE. The total loss L for each image is the weighted sum of L^H , L^D and L^G :

$$L = \sum_{x=1}^{U+V} L^H(h_{fx}, \hat{h}_{fx}) + \alpha \sum_{y=1}^{V} L^D(D_y, \hat{D}_y) + \beta \sum_{z=1}^{U} L^G$$
(11)

where \hat{h}_{fx} and \hat{D}_y denote the predicted joints confidence map and displacements regression map. α and β are constant weight factor to balance three kinds of losses, both set as 0.01. The overall framework of DE is end-to-end trainable via gradient backpropagation.

The overall architecture of DE is illustrated in Figure 2. For an image, DE generates two kinds of feature maps \hat{h}_{Gk} and \hat{h}_{Lk} . through performing NMS and on them, we get predicted joints location of RJG and AJG. GE gives identification tags for RJG and LE provides the connection relation to connect AJG to RGJ. To better present the collaborative work of GE and LE, we add the intermediate illustration. Connected pairs get identification information from GE and joints in GE expand to all joints by connectivity in LE. Based on the displacements and connectivity in LE, MRM refines the low-quality predicted

joints. The final result is generated through the combination of refined results from GE and LE.

4 Experiments

4.1 Experiment setup

Datasets We evaluate the proposed Double Embedding model on three widely used benchmarks for multi-person pose estimation: MSCOCO [21] dataset, MPII [22] dataset and CrowdPose [23] dataset.

The MSCOCO [21] dataset contains over 200, 000 images and 250, 000 person instances labeled with 17 keypoints. COCO is divided into train/val/test-dev sets with 57k, 5k and 20k images respectively. MPII [22] dataset contains 5,602 images of multiple persons. Each person is annotated with 16 body joints. Images are divided into 3,844 for training and 1,758 for testing. MPII also provides over 28,000 annotated single-person pose samples. The CrowdPose[23] dataset consists of 20,000 images, containing about 80,000 person instances. The training, validation and testing subset are split in proportional to 5:1:4. CrowdPose has more crowded scenes than the COCO and MPII, and therefore is more challenging for multi-person pose estimation.

Data augmentation We follow the conventional data augmentation strategies in experiments. For MSCOCO and CrowdPose datasets, we augment training samples with random rotation ($[-30^\circ, 30^\circ]$), random scale ([0.75, 1.5]), random translation ([40, 40]) and random horizontally flip to crop input images to 640x640 with padding. For MPII dataset, random scale is set as ([0.7, 1.3]) while other augmentation parameters are set the same as MSCOCO and Crowd-Pose datasets.

Evaluation metric For COCO and CrowdPose datasets, the standard evaluation metric is based on Object Keypoint Similarity (OKS):

$$OKS = \frac{\sum_{i} exp(-d_{i}^{2}/2s^{2}k_{i}^{2})\delta(v_{i} > 0)}{\sum_{i} \delta(v_{i} > 0)}$$
(12)

where d_i is the Euclidean distance between the predicted joints and ground truth, v_i is the visibility flag of the ground truth, s is the object scale, and k_i is a per-keypoint constant that controls falloff. The standard average precision and recall scores are shown as: $AP^{50}(AP \text{ at OKS} = 0.50), AP^{75}, AP$ (the mean of AP scores at 10 positions, OKS = 0.50, 0.55, ..., 0.90, 0.95; AP^M for medium objects, AP^L for large objects, and AR at OKS = 0.50, 0.55, ..., 0.90, 0.955.

For MPII dataset, the standard evaluation metric is PCKh (head-normalized probability of correct keypoint) score. A joint is correct if it falls within *al* pixels of the groundtruth position, where α is a constant and *l* is the head size that corresponds to 60% of the diagonal length of the ground-truth head bounding box. The PCKh@0.5 ($\alpha = 0.5$) score is reported.

Implementation For COCO dataset, we use standard validation set for ablation studies while use test-dev set to compare with other state-of-the-art methods. or CrowdPose dataset, we use CrowdPose train and val set to train our

model, and use test set for validation. For MPII dataset, following [36], we randomly select 350 groups of multi-person training samples as the validation dataset and use the remaining training samples and all single-person images as train dataset. We use the Adam optimizer [37]. For COCO and CrowdPose datasets, the base learning rate is set to 1e-3, and dropped to 1e-4 and 1e-5 at the 200th and 260th epochs respectively. We train the model for a total of 300 epochs. For MPII dataset, we initialize learning rate by 1e-3. We train the model for 260 epochs and decrease learning rate by a factor of 2 at the 160th, 180th, 210th, 240th epoch. Following [HigherHRNet-30], we adopt flip test for all the experiments.

Method	AP	AP^{50}	AP^{75}	AP^M	AP^L
w/o multi-scale test					
CMU-Pose[17]†	61.8	84.9	67.5	57.1	68.2
RMPE[11]	61.8	83.7	69.8	58.6	67.6
Associate Embedding[18]	62.8	84.6	69.2	57.5	70.6
Mask-RCNN[9] [†]	63.1	87.3	68.7	57.8	71.4
G-RMI[8]†	64.9	85.5	71.3	62.3	70.0
PersonLab[19]	66.5	88.0	72.6	62.4	72.3
PifPaf[33]	66.7	-	-	62.4	72.9
SPM[36]	66.9	88.5	72.9	62.6	73.1
HigherHRNet[38]	68.4	88.2	75.1	64.4	74.2
$CPN[12]\dagger$	72.1	91.4	80.0	68.7	77.2
DoubleEmbedding(Ours)	69.7	88.4	76.9	65.8	75.1
w/ multi-scale test					
Hourglass[18]	65.5	86.8	72.3	60.6	72.6
Associate Embedding[18]	65.5	86.8	72.3	60.6	72.6
PersonLab[19]	68.7	89.0	75.4	64.1	75.5
HigherHRNet[38]	70.5	89.3	77.2	66.6	75.8
DoubleEmbedding(Ours)	71.6	89.5	78.6	68.8	76.0

Table 1. Comparison with state-of-the-arts on COCO2017 test-dev dataset. Top: w/o multi-scale test. Bottom: w/ multi-scale test.

† indicates top-down methods

4.2 Results on COCO dataset

Comparison with state-of-the-arts In Table 1, we compare our proposed model with other state-of-the-arts methods on COCO2017 test-dev dataset. We test the run time of single-scale inference, the proposed method realizes the balance on speed and accuracy. We achieve the competitive accuracy which outperforms most existing bottom-up methods. Compared to the typical top-down

method CPN[12], we narrow the gap to top-down method in accuracy with less complexity. This demonstrates the effectiveness of DE on multi-person pose estimation.

Ablation analysis We conduct ablation analysis on COCO2017[21] validation dataset without multi-scale test. We evaluate the impact of the introduced hierarchical short-range displacements that factorize the basic displacements. Also, the effect of MRM is studied. MRM is implemented based on the hierarchical displacements, so MRM is non-existent without hierarchical displacements. Results are shown in Table 2. which shows that with basic displacements only, DE achieves 69.3% mAP. By introducing hierarchical displacements, performance improves to 70.7% mAP with 1.3% mAP increasing. Based on the hierarchical displacements, MRM further improve 0.9% mAP. The result shows the effectiveness of hierarchical displacements and MRM.

Table 2. Ablation experiments on COCO validation dataset.

Model Settings				Pose Estimation							
Basic Dis.	Hierar Dis.	MRM	AP	AP^{50}	AP^{75}	AP^M	AP^L				
\checkmark			69.3	87.0	75.9	65.1	76.6				
\checkmark	\checkmark		70.7	88.1	76.8	66.5	76.4				
\checkmark	\checkmark	\checkmark	71.6	88.3	77.5	66.9	77.8				

In addition, we analyze the impact of the hyper-parameter τ which decides the regression area for joints confidence map and displacements in Section 3. We observe the performance of proposed model as τ varies from 1 to 20. As shown in Figure 3, the performance monotonically improves as τ increases from 1 to 7. When $7 < \tau < 10$, performance remains unchanged as τ increases. When $\tau > 10$, performance drops as τ increases. This can be explained by the distribution of positive samples of the dataset. When τ increases in the range between 1 and 7, positive samples increase and larger effective area of joints is counted in joints confidence and displacements regression in training. When τ increases between 7 and 10, effective information and background noise increases with equivalent effect. When $\tau > 10$, more background noise is countered as positive samples, regression area of joints overlaps with each other as τ keeps increasing. Smaller τ means less complexity, thus we set $\tau = 7$ for balancing the accuracy and efficiency.

Qualitative results Qualitative results on COCO dataset are shown in the top row of Figure 4. The proposed model performs well in challenging scenarios, e.g., pose deformation (1st and 2nd examples), person overlapping and self-occlusion (3rd example), crowded scene (4nd example), and scale variation and small-scale prediction (5st example). This presents the effectiveness of our method.



Fig. 3. Studies on hyper-parameter τ , which decides the regression area for joints confidence map and displacements.

Table 3. Comparison with state-of-the-arts on the full testing set of MPII dataset.

Method	Head	Sho	Elb	Wri	Hip	Knee	Ank	Total	$\operatorname{Time}[\mathbf{s}]$
lqbal and Gall[16]	58.4	53.9	44.5	35.0	42.2	36.7	31.1	43.1	10
Insafutdinov et al.[15]	78.4	72.5	60.2	51.0	57.2	52.0	45.4	59.5	485
Levinkov et al.[15]	89.8	85.2	71.8	59.6	71.1	63.0	53.5	70.6	-
Insafutdinov et al.[39]	88.8	87.0	75.9	64.9	74.2	68.8	60.5	74.3	-
Cao et al.[32]	91.2	87.6	77.7	66.8	75.4	68.9	61.7	75.6	0.6
Fang et al.[11]	88.4	86.5	78.6	70.4	74.4	73.0	65.8	76.7	0.4
Newell and Deng[13]	92.1	89.3	78.9	69.8	76.2	71.6	64.7	77.5	0.25
Fieraru et al. [40]	91.8	89.5	80.4	69.6	77.3	71.7	65.5	78.0	-
SPM[36]	89.7	87.4	80.4	72.4	76.7	74.9	68.3	78.5	0.058
DoubleEmbedding (Ours)	91.9	89.7	81.6	74.9	79.8	75.8	71.5	80.7	0.21

 ${\bf Table \ 4.} \ {\rm Ablation \ experiments \ on \ MPII \ validation \ dataset.}$

Model Settings			Pose Estimation							
Basic Dis.	Hierar Dis.	MRM	Head	Sho	Elb	Wri	Hip	Knee	Ank	Total
\checkmark			92.1	88.4	78.3	68.9	77.5	73.6	63.7	77.5
\checkmark	\checkmark		92.3	89.2	79.8	71.3	78.1	74.8	66.2	78.8
\checkmark	\checkmark	\checkmark	92.3	90.1	81.2	72.6	79.0	75.7	67.1	79.7

4.3 Results on MPII dataset

Table 3 shows a comparison of the proposed method with state-of-the-arts methods on MPII dataset. The proposed model obtains 80.7% mAP achieving competitive result among other bottom-up methods. In addition, we conduct the ablation study on MPII validation dataset to verify MRM and the hierarchical displacements compared with the basic displacements. As shown in table 4, DE improves from 77.5% mAP to 78.8% mAP by introducing hierarchical displacements. Moreover, on wrists and ankles are significant from 68.9% to 71.3% mAP and 73.6% to 74.8% mAP, respectively. Indicating the effectiveness of hierarchical displacements to factorize the long-range displacements. MRM further improves 1.1% mAP based on hierarchical displacements.

Qualitative results on MPII are shown in the middle row of Figure 4, demonstrating the good performance and robustness of our model in complex scenes such as person scale variations (1st example), large pose deformation (2nd and 3rd examples) and small-scale prediction (3rd example).

4.4 Results on CrowdPose dataset

Method	AP	AP^{50}	AP^{75}	AP^M	AP^L	AP^{H}
Openpose[17]	-	-	-	62.7	48.7	32.3
Mask-RCNN [†] [9]	57.2	83.5	60.3	69.4	57.9	45.8
AlphaPose [†] [11]	61.0	81.3	66.0	71.2	61.4	51.1
$SPPE^{\dagger}[23]$	66.0	84.2	71.5	75.5	66.3	57.4
HigherHRNet[38]	67.6	87.4	72.6	75.8	68.1	58.9
HRNet [†] [41]	71.7	89.8	76.9	79.6	72.7	61.5
DoubleEmbedding (Ours)	68.8	89.7	73.4	76.1	69.5	60.3

Table 5. Comparison with state-of-the-arts on CrowdPose test dataset.

† indicates top-down methods

Table 5 shows experimental results on CrowdPose. The proposed model achieves 68.8% AP which outperforms the existing bottom-up methods. but the performance is still lower than the state-of-the-art top-down method, HRNet which has intrinsic advantage in accuracy due to its processing flow. However, it narrows the accuracy gap between other bottom-up methods and top-down methods with less complexity. The performance on CrowdPose dataset indicates the robustness of our method in crowded scene.

Qualitative results on CrowdPose dataset are shown in the bottom row of Figure 4. The result verifies the effectiveness of our model in complex scenes, e.g., ambiguity pose and small-scale prediction (1st example), self-occluded (2nd example), cluttered background (3rd example) and person overlapping and crowded

scene (4th example).



Fig. 4. Qualitative results on MSCOCO dataset (top), MPII dataset (middle) and CrowdPose dataset (bottom).

5 Conclusion

In this paper, we propose the Double Embedding (DE) method for multi-person pose estimation. Through Global Embedding (GE) and Local Embedding (LE), we achieve parallel implementation of person detection and joints detection, overcoming the intrinsic disadvantages of the conventional two-stage strategy on multi-person pose estimation. GE reduces the person instance detection problem to identifying a group of joints and LE connects and identifies the rest joints hierarchically. Based on LE, we design Mutual Refine Machine (MRM) to further enhance the performance for dealing with complex scenarios. We implement DE based on CNNs with end-to-end learning and inference. Experiments on three main benchmarks demonstrate the effectiveness of our model. DE achieves the competitive results among existing bottom-up methods and narrows the gap to the state-of-the-art top-down methods with less complexity.

References

- 1. Chéron, G., Laptev, I., Schmid, C.: P-cnn: Pose-based cnn features for action recognition. (2015)
- Li, Y., Zhou, S., Huang, X., Xu, L., Ma, Z., Fang, H., Wang, Y., Lu, C.: Transferable interactiveness prior for human-object interaction detection. CoRR abs/1811.08264 (2018)

- 3. Fang, H.S., Cao, J., Tai, Y.W., Lu, C.: Pairwise body-part attention for recognizing human-object interactions. (2018)
- 4. Qian, X., Fu, Y., Xiang, T., Wang, W., Qiu, J., Wu, Y., Jiang, Y.G., Xue, X.: Pose-normalized image generation for person re-identification. (2017)
- Xiao, B., Wu, H., Wei, Y.: Simple baselines for human pose estimation and tracking. (2018)
- Sun, K., Xiao, B., Liu, D., Wang, J.: Deep high-resolution representation learning for human pose estimation. (2019)
- Wang, J., Sun, K., Cheng, T., Jiang, B., Xiao, B.: Deep high-resolution representation learning for visual recognition. IEEE Transactions on Pattern Analysis and Machine Intelligence **PP** (2020) 1–1
- Papandreou, G., Zhu, T., Kanazawa, N., Toshev, A., Tompson, J., Bregler, C., Murphy, K.P.: Towards accurate multi-person pose estimation in the wild. CoRR abs/1701.01779 (2017)
- 9. He, K., Georgia, G., Piotr, D., Ross, G.: Mask r-cnn. IEEE Transactions on Pattern Analysis Machine Intelligence (2018) 1–1
- Huang, S., Gong, M., Tao, D.: A coarse-fine network for keypoint localization. In: 2017 IEEE International Conference on Computer Vision (ICCV). (2017)
- Fang, H., Xie, S., Tai, Y., Lu, C.: Rmpe: Regional multi-person pose estimation. (2016)
- 12. Chen, Y., Wang, Z., Peng, Y., Zhang, Z., Yu, G., Sun, J.: Cascaded pyramid network for multi-person pose estimation. (2017)
- Newell, A., Yang, K., Deng, J.: Stacked hourglass networks for human pose estimation. (2016)
- Pishchulin, L., Insafutdinov, E., Tang, S., Andres, B., Andriluka, M., Gehler, P., Schiele, B.: Deepcut: Joint subset partition and labeling for multi person pose estimation. (2016)
- 15. Insafutdinov, E., Pishchulin, L., Andres, B., Andriluka, M., Schiele, B.: DeeperCut: A deeper, stronger, and faster multi-person pose estimation model. (2016)
- Iqbal, U., Gall, J.: Multi-person pose estimation with local joint-to-person associations. (2016)
- 17. Osokin, D.: Real-time 2d multi-person pose estimation on cpu: Lightweight openpose (2018)
- Newell, A., Deng, J.: Associative embedding: End-to-end learning for joint detection and grouping. CoRR abs/1611.05424 (2016)
- Papandreou, G., Zhu, T., Chen, L., Gidaris, S., Tompson, J., Murphy, K.: Personlab: Person pose estimation and instance segmentation with a bottom-up, partbased, geometric embedding model. CoRR abs/1803.08225 (2018)
- Cheng, B., Xiao, B., Wang, J., Shi, H., Zhang, L.: Higherhrnet: Scale-aware representation learning for bottom-up human pose estimation. In: 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). (2020)
- Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. (2014)
- Andriluka, M., Pishchulin, L., Gehler, P., Schiele, B.: Human pose estimation: New benchmark and state of the art analysis. In: Computer Vision and Pattern Recognition (CVPR). (2014)
- 23. Li, J., Wang, C., Zhu, H., Mao, Y., Fang, H.S., Lu, C.: Crowdpose: Efficient crowded scenes pose estimation and a new benchmark. (2018)
- Ren, S., He, K., Girshick, R., Sun, J.: Faster r-cnn: Towards real-time object detection with region proposal networks. IEEE Transactions on Pattern Analysis and Machine Intelligence (2015)

- 16 Yiming Xu, Jiaxin Li et al.
- 25. Lin, T.Y., Dollár, P., Girshick, R., He, K., Hariharan, B., Belongie, S.: Feature pyramid networks for object detection. (2016)
- Cheng, B., Wei, Y., Shi, H., Feris, R.S., Xiong, J., Huang, T.S.: Revisiting RCNN: on awakening the classification power of faster RCNN. CoRR abs/1803.06799 (2018)
- 27. Cheng, B., Wei, Y., Shi, H., Feris, R., Xiong, J., Huang, T.: Decoupled classification refinement: Hard false positive suppression for object detection. (2018)
- Gkioxari, G., Hariharan, B., Girshick, R., Malik, J.: Using k-poselets for detecting people and localizing their keypoints. In: 2014 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2014)
- Min, S., Savarese, S.: Articulated part-based model for joint object detection and pose estimation. In: IEEE International Conference on Computer Vision, ICCV 2011, Barcelona, Spain, November 6-13, 2011. (2011)
- Wei, S., Ramakrishna, V., Kanade, T., Sheikh, Y.: Convolutional pose machines. In: 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2016) 4724–4732
- Jaderberg, M., Simonyan, K., Zisserman, A., Kavukcuoglu, K.: Spatial transformer networks. (2015)
- Cao, Z., Simon, T., Wei, S., Sheikh, Y.: Realtime multi-person 2d pose estimation using part affinity fields. CoRR abs/1611.08050 (2016)
- Kreiss, S., Bertoni, L., Alahi, A.: Pifpaf: Composite fields for human pose estimation. CoRR abs/1903.06593 (2019)
- Carreira, J., Agrawal, P., Fragkiadaki, K., Malik, J.: Human pose estimation with iterative error feedback. CoRR abs/1507.06550 (2015)
- 35. Girshick, R.: Fast r-cnn. Computer Science (2015)
- Nie, X., Zhang, J., Yan, S., Feng, J.: Single-stage multi-person pose machines (2019)
- Kingma, D., Ba, J.: Adam: A method for stochastic optimization. Computer Science (2014)
- Cheng, B., Xiao, B., Wang, J., Shi, H., Huang, T.S., Zhang, L.: Higherhrnet: Scale-aware representation learning for bottom-up human pose estimation (2019)
- Insafutdinov, E., Andriluka, M., Pishchulin, L., Tang, S., Levinkov, E., Andres, B., Schiele, B.: Articulated multi-person tracking in the wild. CoRR abs/1612.01465 (2016)
- 40. Fieraru, M., Khoreva, A., Pishchulin, L., Schiele, B.: Learning to refine human pose estimation. CoRR abs/1804.07909 (2018)
- 41. Sun, K., Xiao, B., Liu, D., Wang, J.: Deep high-resolution representation learning for human pose estimation. CoRR **abs/1902.09212** (2019)