

Peeking into occluded joints: A novel framework for crowd pose estimation

Lingteng Qiu^{1,2†}, Xuanye Zhang^{1†}, Yanran Li³, Guanbin Li⁴, Xiaojun Wu²,
Zixiang Xiong⁵, Xiaoguang Han^{1*}, and Shuguang Cui¹

¹Shenzhen Research Institute of Big Data, The Chinese University of Hongkong,
Shenzhen²Harbin Institute of Technology(Shenzhen)³Bournemouth University⁴Sun
Yat-sen University⁵Texas A& M University

Abstract. Although occlusion widely exists in nature and remains a fundamental challenge for pose estimation, existing heatmap-based approaches suffer serious degradation on occlusions. Their intrinsic problem is that they directly localize the joints based on visual information; however, the invisible joints are lack of that. In contrast to localization, our framework estimates the invisible joints from an inference perspective by proposing an Image-Guided Progressive GCN module which provides a comprehensive understanding of both image context and pose structure. Moreover, existing benchmarks contain limited occlusions for evaluation. Therefore, we thoroughly pursue this problem and propose a novel OPEC-Net framework together with a new Occluded Pose (OCPose) dataset with 9k annotated images. Extensive quantitative and qualitative evaluations on benchmarks demonstrate that OPEC-Net achieves significant improvements over recent leading works. Notably, our OCPose is the most complex occlusion dataset with respect to average IoU between adjacent instances. Source code and OCPose are publicly available at <https://lingtengqiu.github.io/2020/03/22/OPEC-Net/> soon.

Keywords: Pose Estimation, Occlusion, Progressive GCN

1 Introduction

Human pose estimation is a long-standing problem in Computer Vision. It has still attracted increasing attentions in recent years due to rising demands for wide range of applications which require human pose as input [2, 4, 8, 12, 16, 19, 23]. Despite the significant progress achieved in this area by advanced deep learning techniques [13, 7, 3, 9, 25], pose estimation in crowd scenarios still remains extremely challenging due to the intractable occlusion problem.

Trending models for crowd pose estimation strongly rely on heatmap representation for joints estimation: albeit being effective for visible joints, these methods still suffer performance degradation on occlusions and this is due to

[†]The first two authors should be considered as joint first authors.

* Corresponding Author, hanxiaoguang@cuhk.edu.cn

the fact that, since invisible joints are hidden, it is infeasible to directly localize them. To date, researchers have made painstaking efforts and complicated remedies in developing heatmap models and improving their accuracy of localization. However, the occlusion problem has only received little attention and only few attempts have been made into solving it. As illustrated in Fig 1, the current state-of-the-art work still produces very awkward poses and fails to estimate the occluded joints.




Fig. 1. The current SOTA method [13] (left) VS our method (right). Our method demonstrates a more natural and accurate estimation for occluded joints

Occlusion is an intractable challenge in pose estimation due to the complicated background context, complex intertwined human poses and arbitrary occluded shape. To reveal the hidden joints, it becomes necessary to have a comprehensive inference method rather than simple localization. *Our key insight is that the invisible joints are strongly related to the contextual understanding of the image and structural understanding of the human pose.* For example, humans can easily infer the location of invisible joints using clues derived from the action type and the image context. Therefore, we delve deeper into the clues needed for invisible joints inference and propose a novel framework OPEC-Net to incorporate these clues for multi-person pose estimation. To achieve this goal, two stages are proposed in our framework: Initial Pose Estimation and GCN-based Pose Correction. The first stage generates heatmaps to produce an initial pose and the subsequent correction stage adjusts the initial pose obtained from the heatmaps by an Image-Guided Progressive GCN (IGP-GCN) module.

The correction module deals with the image context and pose structure clues in the following aspects: (1) The human body structure provides the essential constraint information between joints. For this reason, the correction module is designed as a GCN-based network, which offers an explicit way of modelling the body structural information that is advantageous for correcting the joints. (2) Another important clue to infer the invisible joints is their related image context. Considering that, our GCN network is specially designed in an Image-Guided way: The IGP-GCN feeds both the coordinate of joints and also the image features extracted at the location of joints as input to each graph node. Therefore, the multi-scale image features from the heatmap modules are fed into

the IGP-GCN in a progressive way, so that large displacements can be learned steadily. This enables the IGP-GCN to not only capture pose structural information, but also the contextual image information at the same time. (3) In a crowd scenario, human interaction information becomes vital to infer poses. Therefore, we further formulate a CoupleGraph by connecting the corresponding joints of two instances, making the interaction between the pair of people contribute to our estimation results as well. However, the multi-scale image features learnt for heatmap estimation are not compatible for the coordinate correction module. Thus, a Cascaded Feature Adaption (CFA) strategy is introduced to process the features first: since the finer image feature has lost more global contextual information, we fuse the low-level features with high-level features following a cascaded design in order to strengthen their contextual information.

Finally, our framework is trained in an end-to-end fashion and addresses the occlusion problem in an elegant manner. Interestingly, the heatmap module and coordinate GCN module are complementary in our framework: the quantisation error introduced from the heatmap modules can be addressed by the IGP-GCN and, at the same time, the heatmap modules offer a more accurate initial value for IGP-GCN that benefits the correction.

We conduct comprehensive experiments and introduce a new dataset to evaluate our framework. While occlusion cases are ubiquitous in crowded scenarios, only few existing benchmarks include enough complex examples tailored to the evaluation of this problem. Thus, it becomes necessary to have datasets that not only contain light occlusions but also include heavily occluded scenes, such as waltz and wrestling, in which individuals are intertwined in complex ways. However, this field is still lacking such datasets because annotating human poses in heavily occluded scenes is very difficult and requires massive manual work. Therefore, we introduce a new dataset called **Occluded Pose (OCPose)** that includes more complex occlusions. We manually label all the 18k ground truth human poses of the 9k images in OCPose. We also compare the average intersection over union (IoU) using typical datasets. MSCOCO [15] and MPII [1] have less than 5% data with IoU higher than 30%, in contrast, our dataset OCPose contains 90% data with IoU higher than 30%.

In summary, the contributions of this work are:

- To the best of our knowledge, this is the first attempt at tackling the challenging problem of occluded joints from digging the image context and pose structure clues in an inference perspective. A novel framework, named OPEC-Net, is proposed, which significantly outperforms existing methods.
- Our approach designs a novel Image-Guided Progressive GCN to accommodate the structural pose information and contextual image information for correction in a single pass.
- We contribute a carefully annotated 9K human pose dataset **OCPose** that includes highly challenging occluded scenes. To the best of our knowledge, OCPose is one of the datasets that contains the most complex occlusions to date. The OCPose dataset will be released to the public to facilitate research in the pose estimation field.

2 Related Works

Heatmap-based Models for pose estimation.

Models for multi-person pose estimation (MPPE) can be divided into two categories, namely bottom-up and top-down approaches. The bottom-up methods first detect the joints and then assign them to the matching person. Pioneer works of bottom-up methods [22, 11, 3, 18, 27] attempted to design different joint grouping strategies. Newell *et al.* [18] introduced a stacked hourglass network to utilize the tagging heatmap. DeepCut [22] presented an Integer Linear Program (ILP) and Zanfir *et al.* [27] grouped joints by learned scoring functions. Cao *et al.* [3] proposed a novel 2d vector field Part Affinity Fields (PAFs) for association as well. However, these prior works all have a serious deficiency that the invisible joints will decrease the performance drastically.

In the second category, the top-down methods first detect all people in the scene and then perform pose estimation for each person. Most of the existing top-down approaches [9, 20, 7] focused on proposing a more effective human detector to obtain better results. Fang *et al.* [7] proposed a framework which is more robust for the redundant human bounding box. Li *et al.* [13] designed a global maximum joints association algorithm to address the association problem in crowd scenarios. Chen *et al.* [5] presented a Cascaded Pyramid Network (CPN) framework in order to relieve the problem from these occluded key points. Moon *et al.* [17] proposed a human pose refinement network that estimates a refined pose from a tuple of an input image and input pose, using the error statistics as prior information to generate synthetic poses. Nevertheless, all of these strategies are unable to adequately reduce errors, especially in the severe occlusion cases, where one bounding box captures joints of multiple people. Most of the mainstream approaches are heatmap-based and thus are limited to estimating invisible joints which are lack visual information. Therefore, we propose an OPEC-Net which completely differs from these works and is able to estimate invisible joints by inference rather than by localization.

GCN for pose modelling. The human body shows a natural graph structure, so that some advanced work constructed graph networks to address human pose related problems, such as action recognition [26], motion prediction [14], and 3D pose regression [29, 6]. These works intuitively form the natural human pose as a graph and apply convolutional layers on it. Compared to other approaches, Graph Convolutional Networks demonstrate one compelling advantage when dealing with human pose modelling problems: they are more effective in capturing dependency relationships between joints.

Previous work [26, 14] achieved a significant gap of improvements in human motion understanding by forming the spatial and temporal relationships as edges in graphs. Moreover, pose regression from 2D to 3D is a natural graph prediction problem so that a new SemGCN [29] is proposed in this field. However, GCN frameworks are never introduced for key points detection problem such as MPPE. In comparison, our graph network is specially designed for key points detection and contains a progressive learning strategy that is guided by image features.

3 OPEC-Net: Occluded Pose Estimation and Correction

Existing pose estimation approaches achieve striking results on visible joints but produce wildly inaccurate outcomes on invisible ones. This is mainly due to the fact that localizing invisible joints from the heatmap is very challenging since they are occluded and there is a lack of visual information. To rectify this shortcoming, we introduce a novel framework that infers invisible joints from the image contextual and pose structural clues.

Considering that, we generate an initial pose from a heatmap-based module and process it into a GCN-based joints correction module to learn their precise position. An Image-Guided GCN network (IGP-GCN) and a Cascaded Feature Adaption module is proposed in the correction stage. The IGP-GCN network exploits the human body structure and image context together to optimize the estimation results. By learning the displacements in a progressive way, it also offers a stable way to achieve more accurate results.

The heatmap and coordinate modules in our framework are actually inter-dependent. Due to our heatmap inference network, the IGP-GCN module has a more accurate pose initialization, which also contributes to a more precise local contextual understanding, before conducting corrections. On the other hand, the coordinate based IGP-GCN also addresses the limitation of heatmap modules: due to a size limit, heatmap representation usually causes quantisation errors for joints estimations. Our IGP-GCN design tackles this issue by converting the heatmap into coordinate representation. The overall framework and the proposed OPEC-Net module is illustrated in Fig 2.

3.1 Initial Pose Estimation from Heatmap-based modules

In this stage, AlphaPose+ [13] is employed as the base module to generate a heatmap for visible joints. This is a top-down approach, which first detects a bounding box for each person and then performs instance-level human pose estimation. We describe the process for an instance-level human pose in the following.

Firstly, the three layers of the base module’s decoder generate three corresponding feature maps with different levels of fine details: a coarse feature map \mathcal{F}_1 , a middle feature map \mathcal{F}_2 and a fine feature map \mathcal{F}_3 . The base module outputs a heatmap which has high confidence for visible joints. The estimated pose from the heatmap H can be denoted as P , which contains estimation results for each joint:

$$\{ \langle x^1, y^1, c^1 \rangle, \langle x^2, y^2, c^2 \rangle, \dots, \langle x^k, y^k, c^k \rangle \} \quad (1)$$

where x^j and y^j are the position of the j th joint, c^j the confidence score, and k is the number of joints in the skeleton.

3.2 GCN-based Joints Correction

The occluded poses can be inferred easily by humans mostly because of their abundant prior knowledge of implicit body structure and pose properties. More

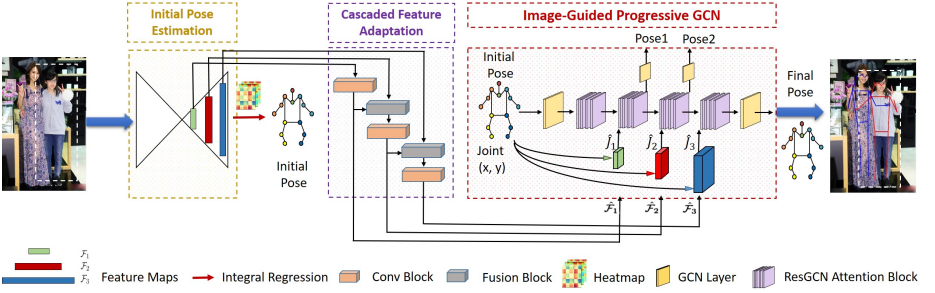


Fig. 2. The schematic diagram of our pipeline. This figure depicts the two stages of estimation for one single pose. The GCN-based pose correction stage contains two modules: the Cascaded Feature Adaptation and the Image-Guided Progressive GCN. Firstly a base module is employed to generate heatmaps. After that, an integral regression method [24] is employed to transform the heatmap representation into a coordinate representation, which can be the initial pose for the GCN network. The initial pose and the three feature maps from the base module are processed in the Image-Guided Progressive GCN. The multi-scale feature maps are updated through the Cascaded Feature Adaptation module and put into each ResGCN Attention blocks. \hat{J}_1 , \hat{J}_2 and \hat{J}_3 are the node features excavated on the related location (x, y) from image features. The error of Initial Pose, Pose1, Pose2, and Final Pose are all considered in the objective function. Then the OPEC-Net is trained entirely to estimate the human pose. The details of the whole framework are described in Section 3

specifically, a natural human pose is highly constrained by the environments and human body property, such as the biomechanical structure of the human body and implications in the environments. In light of that, we propose an Image-Guided graph network for correction which takes the initial pose generated from the above modules and adjusts the estimation results according to the implicit relationship of joints.

Graph Formulation. The human body skeleton has a natural hierarchical graph structure. Previous researches on MPPE merely utilize this information by a primitive graph matching strategy. We claim that the implicit relationships between different joints are helpful to guide position estimation. We thus construct an intuitive graph $G = (V, E)$ to formulate the human pose with N joints. V is the node set in G which can be denoted as $V = \{v_i | i = 1, 2, \dots, N\}$. $E = \{v_i v_j | \text{if } i \text{ and } j \text{ are connected in the human body}\}$ is the edge set which refers to the human body limbs. The adjacent matrix of G refers to matrix $A = \{a_{ij}\}$, with $a_{ij} = 1$ when v_i and v_j are neighbours in G or $i = j$, otherwise $a_{ij} = 0$.

For every node, the input feature G_i^j is the joint estimation result $\langle x_i^j, y_i^j, c_i^j \rangle$, where i is i th pose and j is the j th joint of the skeleton. We denote $G_i \in \mathbb{R}^{L \times N}$ as the input feature of the i th pose in the training set, where L is the feature dimension.

Image-Guided Progressive GCN Network. First of all, to bridge the gap that the gradient cannot be backpropagated between the heatmap produced by basic module and joint position, we utilize the soft-argmax[24] to get \hat{J}_i^k , where \hat{J}_i^k is the position estimation of the k th joint. The core methodology proposed in our work is the Image-Guided Progressive GCN for Correction. In this network, the image context and pose structure clues for invisible joints inference are merged together in an innovative way. The details of each layers and ResGCN Attention Blocks are describe in supplementary materials.

(1) The estimated position of invisible joints from the base module is sometimes far from their correct locations and this makes it challenging to directly regress their displacements. Therefore, we design an intuitive coarse-to-fine learning mechanism in the coordinate-based module, which builds a ***progressive GCN architecture*** and leverages the performance steadily by enforcing multi-scale image features in a progressive manner.

(2) The coordinate-based module lacks ***local context information***. Consequently, we excavate the related image features for each joints position and fuse them into the module. In other words, we improve the pose estimation results by incorporating image feature maps $\hat{\mathcal{F}}_1, \hat{\mathcal{F}}_2$, and $\hat{\mathcal{F}}_3$. Specifically, we design cascaded ResGCN attention blocks to grasp the useful information that is stored in the feature maps but lost in initial pose \hat{P}_i . The three feature maps are ordered from coarse to fine according to their size of receptive fields. After that, we employ a grid sample method that obtains the j th joint feature by excavating the feature located in $\langle x_i^j, y_i^j \rangle$ on the related coordinate weight feature map. Every pose leads to three node feature vectors \hat{J}_1, \hat{J}_2 , and \hat{J}_3 extracted following this process. Finally, these node features are fed into the ResGCN attention blocks accordingly.

Cascaded Feature Adaption (CFA). Feature maps $\mathcal{F}_1, \mathcal{F}_2$, and \mathcal{F}_3 should be adaptive to provide more effective information to the IGP-GCN. Moreover, the low-level feature and high-level feature are fused in the cascaded design in order to enlarge their respective fields resulting in the updated feature as being more informative. The details of the Conv Blocks and Fusion Blocks are present in the supplementary materials.

3.3 Loss Functions

The objective function of our OPEC-Net module can now be formulated. We denote the training set as Ω , the ground truth pose in Ω as P_i , and the output pose of j th ResGCN Attention block as \hat{P}_{ij} . From heatmap representation to coordinate representation, the integral regression method produces an initial pose \hat{P}_{init} . Hence, the total loss is defined as the sum of the rectified loss of poses from IGP-GCN and initial loss of initial pose:

$$Loss = \min_{\theta} \sum_{i \in \Omega} \left(\sum_{j=1}^n \lambda_j |(\hat{P}_{ij} - P_i) \odot M| + |(\hat{P}_{init} - P_i) \odot M| \right) \quad (2)$$

The term $|\hat{P}_{ij} - P_i|$ indicates the calculation of the L_1 loss between our estimated pose and the ground truth, n is the number of ResGCN attention blocks in the model. In this work, we set $n = 3$. We sum up all the errors of the produced pose from each block and assign a parameter λ_j to control the weight. All the trainable parameters in our network are denoted as θ . $M \in \mathbb{Z}_2^N$ is a binary mask where the element in M corresponds to 1 when the related joint has a ground truth label, otherwise it is 0. The \odot denotes the element-wise product operation so that we only take into account the errors on the joints with ground truth.

4 Our Occluded Pose Dataset

We build a new dataset, called **Occluded Pose** (OCPose), that includes more heavy occlusions to evaluate the MPPE. It contains challenging invisible joints and complex intertwined human poses. We mostly consider the couple pose scenes, such as dancing, skating, and wrestling, because they have more reliable annotations and practical utility. This section gives details of data collection, data annotation, and data statistics.

Data Collection. The ground truth of the human pose can be hard to recognize when the occlusions are extremely heavy. Thus, we majorly collect videos of two-person interactions since they are much easier to annotate and because the volunteer can infer the pose from contextual information. We first search for videos from the Internet by using keywords such as boxing, dancing, and wrestling. We then capture the distinctive images which contain diverse poses and humans from these videos by restricting the interval to be at least 3 seconds. Finally, we manually sift through the clips to select high-quality images. All the images are collected under the permission of privacy issues.

Data Annotation. We develop an annotation tool for the user to bound the area of the couple and then locate two template skeletons to their right positions. Six volunteers are recruited for manual labelling. Each skeleton has 12 joints and the left and right components are distinct. In addition to annotating the bounding box and the human body poses, the volunteers also need to indicate whether the joint is visible or not. To ensure accuracy, we use cross annotation for every image.

Data Statistics. In total, our dataset contains 9000 images and 18000 fully annotated persons. For the training process, the training dataset consists of 5000 images, whereas validation and test dataset each contains 2000 images.

To compare the occlusion level, we evaluate the average intersection over union (IoU) of bounding box on the other public benchmarks, such as CrowdPose [13], OCHuman [28], MSCOCO [15] and MPII[1]. We report the comparison result of these benchmarks in Table 1, which illustrates that our dataset beats down all the other benchmarks on the occlusion level.

Other Dataset. In our approach, we carried out extensive experiments on public benchmarks. Following the typical training procedure, we evaluate the OPEC-Net on our OCPose, CrowdPose [13], MSCOCO [15] and particular oc-

Table 1. The comparison of occlusion level. We count the number of images of each dataset with different levels of occlusion. As shown above, MSCOCO and MPII almost have no heavy occlusions. OCHuman is the state-of-the-art dataset for occlusions but our dataset is larger and contains more severe occlusions

Dataset	Total	IoU>0.3	IoU>0.5	IoU>0.75	Average
CrowdPose	20000	8706 (44%)	2909 (15%)	309 (2%)	0.27
MSCOCO	118287	6504 (5%)	1209 (1%)	106 (<1%)	0.06
MPII	24987	0	0	0	0.11
OCHuman	4773	3264(68 %)	3244(68%)	1082(23%)	0.46
Ours	9000	8105 (90%)	6843 (76%)	2442 (27%)	0.47

cluded dataset OCHuman [28]. CrowdPose dataset is split in a ratio of 5 : 4 : 1 for training, testing and validation respectively. We regard the validation set of OCHuman with 2500 images as our training dataset, and the rest 2273 images for testing. Then we follow the typical training strategy on MSCOCO.

5 Experiments

In this section, extensive quantitative and qualitative experiments are demonstrated to evaluate the effectiveness of our OPEC-Net. Comprehensive ablation studies are carried out to validate the effectiveness of each components.

5.1 Experiments Settings

Implementation Details. For training, we set the parameters $\lambda_1 = 0.3$, $\lambda_2 = 0.5$, $\lambda_3 = 1$ and $epochs = 30$. We feed 10 images in a batch to train the whole framework. The initial learning rate is set to $1e^{-3}$ and decays in a cosine way. The input image size are 384×288 for MSCOCO and 320×256 for the other datasets. An AdamOptimizer is employed to optimize the parameters by backpropagation. For a fair comparison, we filter the proposal of the instances in the background and only focus on the Object Keypoint Similarity (OKS) of targets when we evaluate baselines on our dataset. We implement our model in PyTorch [21] and conduct experiments on one Nvidia GeForce GTX 1080 Ti with 11GB memory. More details are described in the supplementary materials.

Evaluation Metric. We follow the standard evaluation metric of MSCOCO, which is widely used by existing works as well [7, 13, 10, 3]. Specifically, we report the mean Average Precision (mAP) value at 0.5:0.95, 0.5, 0.75, 0.80 and 0.90. In order to grasp the qualified poses for the OPEC-Net training procedure, two rules are formulated to select the proposal. The proposal poses must contain more than 5 visible points and OKS value more than 0.3 to ensure the quality. To enrich the dataset, we also flip the images as a data augmentation strategy. Furthermore, we provide the visualization results of pose estimation.

Baselines. For comparison, we assess the performance with our OPEC-Net module using the three state-of-the-art approaches for MPPE: Mask RCNN [9], AlphaPose+¹ [13] and SimplePose [25]. For a fair comparison, we quote the results of Mask RCNN and SimplePose directly from paper [13] and re-train AlphaPose+ from their public code. For the evaluation on OCPose, CrowdPose and OCHuman, we take AlphaPose+ for the initial pose estimation stage with ResNet-101 as the backbone and Yolo V3 as the detector. For the MSCOCO dataset, we take the public code of SimplePose² for the first stage because it has a higher performance than AlphaPose+ on MSCOCO. Mask RCNN is used as the detector and ResNet-152 is used as the backbone on MSCOCO. OPEC-Net here denotes the framework with a single person as graph, and CoupleGraph denotes the baseline that performs a CoupleGraph based framework after OPEC-Net. All of details about CoupleGraph is illustrated in supplementary.



Fig. 3. The qualitative evaluation of CoupleGraph and OPEC-Net. The left images are generated from OPEC-Net and the right ones come from CoupleGraph

Table 2. The comparison on our OCPose dataset

Method	mAP@0.5:0.95	AP^{50}	AP^{75}	AP^{80}	AP^{90}
Mask RCNN [9]	21.5	49.8	15.9	7.7	0.1
Simple Pose [25]	27.1	54.3	24.2	16.8	4.7
AlphaPose+ [13]	30.8	58.4	28.5	22.4	8.2
OPEC-Net	32.8(+2.0)	60.5	31.1	24.0	9.2
CoupleGraph	33.6(+2.8)	60.8	32.5	25.0	9.8

5.2 Performance Comparison on our OCPose dataset

Quantitative Comparison. The quantitative results are presented in Table 2. Our approach attains the best mAP compared to all the baselines with a considerable margin. OPEC-Net achieves a significant gain which is a surprising

¹ <https://github.com/MVIG-SJTU/AlphaPose/tree/pytorch>.

² <https://github.com/leoxiaobin/pose.pytorch>.



Fig. 4. The results on OCPose, OCHuman and CrowdPose. These are the qualitative comparison results of AlphaPose+ method and OPEC-Net on our datasets. The left pose is estimated by AlphaPose+ method and the right one is ours. The first row is OCPose and the second row represents OCHuman, the rest represents CrowdPose

2.0 mAP@0.5:0.95 improvement compared to the AlphaPose+. Despite of that, a significant 1.0 AP^{90} improvement has been achieved which proves that our OPEC-Net has great ability of inference especially for high level occlusions compared to localization methods. In conclusion, these results validate the prominent effectiveness of our OPEC-Net module on MPPE tasks.

Qualitative Comparison. As illustrated in the first row of Fig 4, our OPEC-Net is capable of correcting the wrong link between joints and estimating the occluded joints while maintaining high performance on visible joints. We make these observations from the results: (1) For the first sample, a superior pose estimation result is provided by our method. Even an error with large displacement can be corrected by OPEC-Net. (2) Moreover, although the second case has difficult sunlight interference, our approach can also adjust the joints to their correct locations. (3) The third group also shows an evidence that our OPEC-Net module produces more natural poses that conform to human body constraints. (4) The fourth figure shows that our method can find the correct link between joints.

CoupleGraph. The evaluations of CoupleGraph are given in Tab 2 and Fig 3. Compared to OPEC-Net, the CoupleGraph baseline also shows an advanced lifting 0.8 mAP@0.5:0.95, which validates the human interaction clues that are quite prominent. As illustrated in Fig 3 CoupleGraph outperforms OPEC-Net significantly in quality. In these human interactive scenarios, the poses estimated by CoupleGraph are more concordant and superior.

5.3 Comparison against State-of-the-Arts on other benchmarks

Extensive evaluations on heavily benchmarked dataset demonstrate the effectiveness of our model for the occlusion problem. The experimental results on existing benchmarks are presented in Table 3, 4 and Fig 4. Our model surpasses all the baselines by a considerable margin.

Table 3. The qualitative result on occlusion dataset

Dataset Method	OCHuman [28]					CrowdPose [13]				
	mAP@0.5:0.95	AP^{50}	AP^{75}	AP^{80}	AP^{90}	mAP@0.5:0.95	AP^{50}	AP^{75}	AP^{80}	AP^{90}
Mask RCNN [9]	20.2	33.2	24.5	18.3	2.1	57.2	83.5	60.3	-	-
SimplePose [25]	24.1	37.4	26.8	22.6	4.5	60.8	81.4	65.7	-	-
AlphaPose+ [13]	27.5	40.8	29.9	24.8	9.5	68.5	86.7	73.2	66.9	45.9
OPEC-Net	29.1(+1.6)	41.3	31.4	27.0	12.8(+3.3)	70.6(+2.1)	86.8	75.6	70.1	48.8(+2.9)

OCHuman As OCHuman is a new benchmark proposed mainly for pose segmentation, we are the first to report all the baseline results on this challenging occlusion dataset. Compared to the AlphaPose+, we achieve maximum 3.3 improvements on AP^{90} . This further validates that our OPEC-Net model is robust even for highly challenging occlusion scenarios.

CrowdPose As shown in Table 3, OPEC-Net drastically lifts 2.1 mAP@0.5:0.95 of the estimation result over AlphaPose+. It is also worth noting that the improvements remain high when the comparison AP terms are high. For example, our model achieves 0.1, 2.4, 3.2 and 2.9 on AP 50, 75, 80 and 90 respectively.

MSCOCO We also present our results on the largest benchmark MSCOCO. Our model only contributes slightly on accuracy improvements. The reason is the key difference between MSCOCO and other datasets – it contains too few occlusion scenarios, especially the severe ones. Moreover, a lot of invisible joints lack annotations on MSCOCO.

Invisible vs. Visible To investigate the effectiveness on invisible (Inv) and visible (V) joints separately, we report the statistics of each type of joints according to the similar rule of OKS. From Tab 5, our OPEC-Net improves mostly on the invisible joints rather than visible joints. In terms of Inv@75, our framework achieves a considerable marginal 3.3% and 4.9% gains on CrowdPose and OC-Pose respectively. On the contrary, the OPEC-Net only improves a maximal 1% on visible joints because our main focus is the invisible joints. This comparison

Table 4. MSCOCO 2017 test-dev set [15]

Method	mAP@0.5:0.95	mAP@0.5	mAP@0.75
AlphaPose+ [13]	72.2	90.1	79.3
Simple Pose [25]	73.7	91.9	81.8
OPEC-Net	73.9(+0.2)	91.9	82.2

also explains why the gains are smaller on MSCOCO datasets than the other datasets that contains more occlusions.

Table 5. Results for Visible and Invisible Joints on CrowdPose and OCPose

Datasets	CrowdPose				OCPose			
	Inv@75	Inv@90	V@75	V@90	Inv@75	Inv@90	V@75	V@90
AlphaPose+	76.2%	57.2%	89.5%	67.8%	50.7%	17.7%	85.2%	55.3%
OPEC-Net	79.5%	58.4%	90.0%	67.8%	55.6%	20.5%	86.2%	55.1%

5.4 Ablation studies

To analyze our model in details, we conduct comprehensive ablative experiments to evaluate the capability of each component and clues we claimed. As illustrated in Table 6, we present the baselines to investigate the impact of each component.

Firstly, to compare the results of other refinement methods, we used cascaded pyramid networks to refine the heatmap of AlphaPose+. For Inv@75, it reaches 0.2 and 0.3 mAP gains on CrowdPose and OCPose, while ours reaches 3.3 and 4.9. For V@75, it reaches 1.2 and 2.3 while ours gains 0.5 and 1.0. As shown, conv-refinement is indeed able to correct visible joints, but it is still challenging for estimating invisible ones. We performed the method by directly using two 256-dim FC layers to regress displacements as well. However, this drops the mAP of AlphaPose+ by 0.6, 1.1, 0.6 on CrowdPose, OCHuman, and OCPose. This demonstrates the necessity of explicit pose structural modelling. Also, We find that directly applying soft-argmax on AlphaPose+ drops the mAP by 0.4, 2.0 and 0.8 on CrowdPose, OCHuman and OCPose. One possible reason is that soft-argmax tends to produce very high confidence on the visible joints but ignore the focus on invisible points.

Meanwhile, we investigate the impact of the image guided strategy that blends the image context with GCN. From (a), a clear decrease around 2.0 mAP@0.5:0.95 is observed, which points out the importance of the Image-Guide strategy. Without the Image-Guided part, a single GCN network improves the performance poorly. This evidence validates that the GCN module must learn under the guidance of image features.

We further investigate each design of the IGP-GCN. From (b) and (c), we can conclude that the strategy of progressive and coarse to fine feature learning is effective. Moreover, the proposed Cascaded Feature Adaption module is analysed

as well. In Table 6, the mAP value of three datasets falls down significantly, demonstrating that the CFA module plays an indispensable role in the whole framework. We remove the Fusion Blocks and report the results in (e), which further proves the effectiveness of the fusion part in the CFA module. We can conclude that the image guidance is the most imperative in the framework. The CFA module brings an average 0.7 mAP@0.5:0.95 gain on three datasets, which manifests the necessity to make image features adaptive for coordinate modules. Overall, these ablation studies overwhelmingly validate that every component is effective and that the clues are informative for invisible joints inference.

Table 6. Ablation study of our OPEC-Net framework (mAP@0.5:0.95)

The evaluation of removal each component	OCHuman	OCPose	CrowdPose
The AlphaPose+ baseline	30.8	27.5	68.5
(a)Without the Image-Guided strategy in GCN	30.8	27.7	68.6
(b)Without the Progressive design in the GCN	32.2	28.3	69.3
(c)Use one feature F_3 instead of multi-scale features	32.5	28.5	69.7
(d)Remove the Cascaded Feature Adaption module	32.1	28.4	69.9
(e)Remove the Fusion Block in CFA module	32.4	28.7	69.6
The full OPEC-Net	32.8	29.1	70.6

6 Conclusion

In this paper, we proposed a novel OPEC-Net module and a challenging Occluded Pose (OCPose) dataset to address the occlusion problem in Crowd Pose Estimation. Two elaborate components, Image-Guided Progressive GCN and Cascaded Feature Adaptation, are designed to exploit the natural human body constraints and image context. We conduct thorough experiments on four benchmarks and ablation studies to demonstrate the effectiveness and provide a variety of insights. The heatmap and coordinate module are proved to work cooperatively and achieve remarkable improvements in all aspects. By making this dataset available, we hope to arouse the attention and increase interest in the investigation of the occlusion problem in pose estimation.

Acknowledgment

The work was supported in part by the Key Area RD Program of Guangdong Province with grant No. 2018B030338001, by the National Key RD Program of China with grant No. 2018YFB1800800, by Natural Science Foundation of China with grant NSFC-61629101 and 61902334, by Guangdong Research Project No. 2017ZT07X152, and by Shenzhen Key Lab Fund No.ZDSYS201707251409055.

References

1. Andriluka, M., Pishchulin, L., Gehler, P., Schiele, B.: 2d human pose estimation: New benchmark and state of the art analysis. In: *Proceedings of the IEEE Conference on computer Vision and Pattern Recognition*. pp. 3686–3693 (2014)
2. Bansal, A., Ma, S., Ramanan, D., Sheikh, Y.: Recycle-gan: Unsupervised video retargeting. In: *Proceedings of the European Conference on Computer Vision (ECCV)*. pp. 119–135 (2018)
3. Cao, Z., Simon, T., Wei, S.E., Sheikh, Y.: Realtime multi-person 2d pose estimation using part affinity fields. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 7291–7299 (2017)
4. Chan, C., Ginosar, S., Zhou, T., Efros, A.A.: Everybody dance now. *arXiv preprint arXiv:1808.07371* (2018)
5. Chen, Y., Wang, Z., Peng, Y., Zhang, Z., Yu, G., Sun, J.: Cascaded pyramid network for multi-person pose estimation. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 7103–7112 (2018)
6. Ci, H., Wang, C., Ma, X., Wang, Y.: Optimizing network structures for 3d human pose estimation. *ICCV* (2019)
7. Fang, H., Xie, S., Tai, Y.W., Lu, C.: Rmpe: Regional multi-person pose estimation. In: *The IEEE International Conference on Computer Vision (ICCV)*. vol. 2 (2017)
8. Gui, L.Y., Zhang, K., Wang, Y.X., Liang, X., Moura, J.M., Veloso, M.: Teaching robots to predict human motion. In: *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. pp. 562–567. IEEE (2018)
9. He, K., Gkioxari, G., Dollár, P., Girshick, R.: Mask r-cnn. In: *Computer Vision (ICCV), 2017 IEEE International Conference on*. pp. 2980–2988. IEEE (2017)
10. Huang, Z., Huang, L., Gong, Y., Huang, C., Wang, X.: Mask scoring r-cnn. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 6409–6418 (2019)
11. Insafutdinov, E., Pishchulin, L., Andres, B., Andriluka, M., Schiele, B.: Deeppercut: A deeper, stronger, and faster multi-person pose estimation model. In: *European Conference on Computer Vision*. pp. 34–50. Springer (2016)
12. Joo, H., Simon, T., Sheikh, Y.: Total capture: A 3d deformation model for tracking faces, hands, and bodies. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 8320–8329 (2018)
13. Li, J., Wang, C., Zhu, H., Mao, Y., Fang, H.S., Lu, C.: Crowdpose: Efficient crowded scenes pose estimation and a new benchmark. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 10863–10872 (2019)
14. Li, M., Chen, S., Chen, X., Zhang, Y., Wang, Y., Tian, Q.: Actional-structural graph convolutional networks for skeleton-based action recognition. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 3595–3603 (2019)
15. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: *European conference on computer vision*. pp. 740–755. Springer (2014)
16. Mehta, D., Sridhar, S., Sotnychenko, O., Rhodin, H., Shafiei, M., Seidel, H.P., Xu, W., Casas, D., Theobalt, C.: Vnect: Real-time 3d human pose estimation with a single rgb camera. *ACM Transactions on Graphics (TOG)* **36**(4), 44 (2017)
17. Moon, G., Chang, J.Y., Lee, K.M.: Posefix: Model-agnostic general human pose refinement network. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 7773–7781 (2019)

18. Newell, A., Huang, Z., Deng, J.: Associative embedding: End-to-end learning for joint detection and grouping. In: *Advances in Neural Information Processing Systems*. pp. 2277–2287 (2017)
19. Panteleris, P., Oikonomidis, I., Argyros, A.: Using a single rgb frame for real time 3d hand pose estimation in the wild. In: *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*. pp. 436–445. IEEE (2018)
20. Papandreou, G., Zhu, T., Kanazawa, N., Toshev, A., Tompson, J., Bregler, C., Murphy, K.: Towards accurate multi-person pose estimation in the wild. In: *CVPR*. vol. 3, p. 6 (2017)
21. Paszke, A., Gross, S., Chintala, S., Chanan, G., Yang, E., DeVito, Z., Lin, Z., Desmaison, A., Antiga, L., Lerer, A.: Automatic differentiation in pytorch (2017)
22. Pishchulin, L., Insafutdinov, E., Tang, S., Andres, B., Andriluka, M., Gehler, P.V., Schiele, B.: Deepcut: Joint subset partition and labeling for multi person pose estimation. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 4929–4937 (2016)
23. 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. In: *Proceedings of the European Conference on Computer Vision (ECCV)*. pp. 650–667 (2018)
24. Sun, X., Xiao, B., Wei, F., Liang, S., Wei, Y.: Integral human pose regression. In: *Proceedings of the European Conference on Computer Vision (ECCV)*. pp. 529–545 (2018)
25. Xiao, B., Wu, H., Wei, Y.: Simple baselines for human pose estimation and tracking. In: *Proceedings of the European Conference on Computer Vision (ECCV)*. pp. 466–481 (2018)
26. Yan, S., Xiong, Y., Lin, D.: Spatial temporal graph convolutional networks for skeleton-based action recognition. In: *Thirty-Second AAAI Conference on Artificial Intelligence* (2018)
27. Zanfir, A., Marinoiu, E., Zanfir, M., Popa, A.I., Sminchisescu, C.: Deep network for the integrated 3d sensing of multiple people in natural images. In: *Advances in Neural Information Processing Systems*. pp. 8410–8419 (2018)
28. Zhang, S.H., Li, R., Dong, X., Rosin, P., Cai, Z., Han, X., Yang, D., Huang, H., Hu, S.M.: Pose2seg: Detection free human instance segmentation. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 889–898 (2019)
29. Zhao, L., Peng, X., Tian, Y., Kapadia, M., Metaxas, D.N.: Semantic graph convolutional networks for 3d human pose regression. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 3425–3435 (2019)