

MULTI-SCALE SUPERVISED NETWORK FOR HUMAN POSE ESTIMATION

Lipeng Ke¹, Ming-Ching Chang², Honggang Qi¹, Siwei Lyu²

¹ University of Chinese Academy of Sciences, Beijing, China

² University at Albany, State University of New York, NY, USA

ABSTRACT

Human pose estimation is an important topic in computer vision with many applications including gesture and activity recognition. However, pose estimation from image is challenging due to appearance variations, occlusions, clutter background, and complex activities. To alleviate these problems, we develop a robust pose estimation method based on the recent deep conv-deconv modules with two improvements: (1) multi-scale supervision of body keypoints, and (2) a global regression to improve structural consistency of keypoints. We refine keypoint detection heatmaps using layer-wise multi-scale supervision to better capture local contexts. Pose inference via keypoint association is optimized globally using a regression network at the end. Our method can effectively disambiguate keypoint matches in close proximity including the mismatch of left-right body parts, and better infer occluded parts. Experimental results show that our method achieves competitive performance among state-of-the-art methods on the MPII and FLIC datasets.

Index Terms— human pose estimation, conv-deconv module, multi-scale supervision, regression network.

1. INTRODUCTION

Human pose estimation refers to the task of estimating body keypoint locations (wrists, elbows, knees, ankles, *etc.*) from images. This task can be very challenging due to the large variability of human body appearances, posture structures, the action being performed, viewing angles, occlusions, and complex backgrounds and lighting conditions; see Fig. 1. Further sophistication of the inference is required when the cases extend to multi-person scenarios.

Human pose estimation has been studied extensively [1]. Classic methods including the use of histogram of oriented gradients (HOG) and deformable parts model (DPM) rely on hand-craft features [2, 3, 4, 5, 6]. With the prosperity of Deep Neural Networks (DNN), Convolutional Neural Networks (CNN) have demonstrated remarkable performance boost in human pose estimation [7, 8, 9, 10, 11]. Tompson *et al.* [12] adopted the *heatmap* representation of body keypoints to improve their localization during training. A Markov random field (MRF) inspired spatial model is used to estimate keypoint relationship. Chu *et al.* [13] propose a transform kernel



Fig. 1. Examples of our human pose estimation on the MPII dataset. Our method can handle complex appearance, view variations, and diverse activities with heavy occlusions.

method to learn local keypoint relationships, which is solved using a bi-directional tree.

Recently, Wei *et al.* [11] use very deep sequential conv-deconv architecture with large receptive fields to directly perform pose matching on the heatmaps. They also enforce *intermediate supervision* between conv-deconv pairs to prevent gradient vanish. The hourglass module proposed by Newell *et al.* [14] is an extension of Wei *et al.* with the addition of residual connections between the conv-deconv sub-modules. The *hourglass* module can effectively capture and combine features across scales. Chu *et al.* [15] adopt stacked hourglass networks to generate attention maps from features at multiple resolutions with various semantics. Yang *et al.* [16] design a Pyramid Residual Module (PRM) to enhance the deep CNN invariance across scales, by learning the convolutional filters on various feature scales.

State-of-the-art DNNs for pose estimation are still limited in the capability of modeling human body structural priors for effective keypoint matching. Existing methods rely on a brute-force approach by increasing network depth to implicitly enrich the keypoint relationship modeling capability. A major weakness in this regard is the ambiguities arising from the occlusions, clutter backgrounds, or multiple body parts in the scene. In the MPII pose benchmark [17], many methods [10, 11, 14, 15, 16] rely on repeating their pose estimation

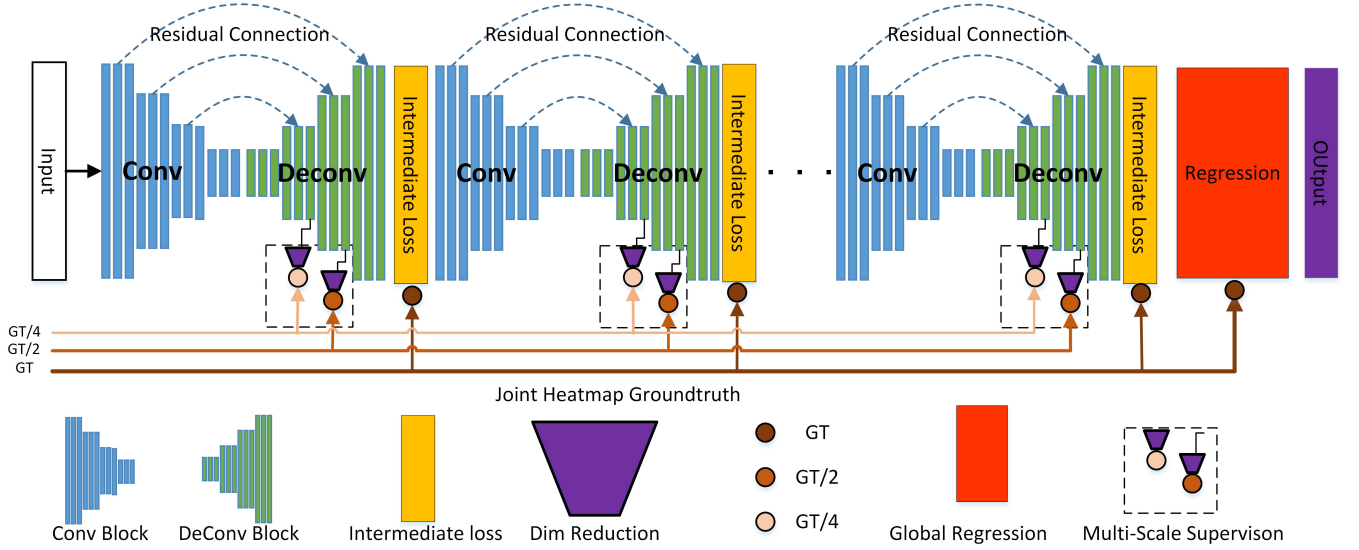


Fig. 2. Our network model consists of four components — (i) conv-deconv modules (blue and green, respectively), (ii) multi-scale supervisions (brown circles next to deconv layers), (iii) intermediate supervision layers (yellow), and (iv) global keypoint regression layers (red).

pipeline multiple times in various scales, in order to improve performance by a small margin using averaging of results. This indicates the lack of an effective solution to handle scale and structural priors in the modeling.

We propose a multi-scale supervised network model consisting of four components depicted in Fig. 2. Our main novelty is two-fold. First, we extend the intermediate supervision to explicitly cover multiple scales at the deconv layers during training. This improves the capability to extract more consistent and representative features across all scales. Our method can then effectively optimize feature representation across scales, because direct supervision is enforced at each scale during learning. Secondly, we use a regression network after the conv-deconv stacks to learn structural priors jointly from the keypoint feature maps from the conv-deconv stacks. This can effectively improve global pose estimation, when compared to existing methods [11, 14, 15, 18] which treat keypoint feature maps independently.

2. MULTI-SCALE SUPERVISED NETWORK MODEL

The proposed multi-scale supervised network is motivated by two key observations. First, in the existing works based on conv-deconv networks [11, 14, 15, 18], accurate body keypoint correspondence depends largely on the consistency of the matching across multiple scales. This leads us to the design of multi-scale supervisions in training our network. Secondly, since each body keypoint heatmap (corresponding to location likelihood) is estimated independently during the conv-deconv steps, structural relationship between individual keypoints are not modeled in the conv-deconv modules. To this end, we apply a global regression network at the end to model the keypoint relationship on top of the heatmaps. This

improves the consistency of body structure in pose estimation in various scenarios: (i) to avoid left-right mismatches, *e.g.* matching a right arm to a left shoulder, (ii) to better handle occlusions, and (iii) to deal with multiple body parts and multiple people in the view.

2.1. Multi-Scale Supervision

We propose to enforce multiple supervision steps at individual deconv layers (shown in Fig. 2) to learn richer multi-scale features for better keypoint localization. As the depth of hour-glass stacks increases, *gradient vanishing* becomes a critical issue during training. *Intermediate supervision* [11] (yellow layers in Fig. 2) between two conv-deconv stacks is a common practice, which by itself can address the gradient vanishing issue to some extent. However intermediate supervision at the original groundtruth scale dose not provide a consistent solution to cohesively supervise feature training across all conv-deconv scales. Our solution is then to apply supervisions to multiple scales of the deconv layers as shown in Fig. 2.

Our multi-scale supervision is an extension of the original intermediate supervision [11]. However our implementation to adopt the multi-scale design is different. Our multi-scale supervision is performed by calculating the residual in each scale regarding the down-sampled groundtruth heatmaps (denoted as GT/8, GT/4, GT/2) at each deconv layer in Fig. 2. Specifically, to make consistent the feature map channels for the computation of keypoint groundtruth heatmap residuals at each scale, we use an 1-by-1 convolutional kernel (purple trapezoid in Fig. 2) to convert the high-dimensional deconv feature maps into individual heatmap for each keypoint. This way, the dimension-reduced feature maps can be directly su-

pervised against the respective scaled groundtruth using mean square error (MSE). We observe that our multi-scale supervision approach can improve the accuracy of keypoint heatmaps (with more focused distributions at keypoints) for use in the next deconv layer and subsequent networks.¹

We describe our multi-scale intermediate loss terms *w.r.t.* the heatmaps of all keypoints as $L2$ loss in the following. For the detection of $N(=16)$ keypoints (head, neck, pelvis, thorax, shoulders, elbows, wrists, knees, ankles, and hips), N heatmaps will be generated after each conv-deconv stack. The loss L_i at the i -th scale compares the predicted heatmaps (of all keypoints) against the ground-truth heatmaps:

$$L_i = \frac{1}{N} \sum_{n=1}^N \sum_{x,y} \|P_n(x,y) - G_n(x,y)\|_2, \quad (1)$$

where $P_n(x,y)$ and $G_n(x,y)$ denote the predicted and the groundtruth heatmaps at the pixel location (x,y) for the n -th keypoint, respectively. The total loss function is the summation across scales, $L = \sum_i L_i$, which is a combination of both the intermediate and multi-scale supervisions.

2.2. Global Keypoint Regression

We use a *fully convolutional* regression network after the conv-deconv stacks to globally refine the multi-scale keypoint heatmaps to improve the pose structural consistency. Our intuition is that the relative positions of arms and legs *w.r.t.* the head/torso represent useful action priors, which can be learned from the regression network by considering feature maps across all scales for pose refinement. Our conv-deconv stacks extract heatmaps which are typically non-Gaussian according to the person’s gesture/activity (as shown in Fig. 3). The regression network then takes the multi-scale heatmaps as input, and match to the input image at respective scales. This way the regression network can effectively oversee the heatmaps across all scales for fine-tuning.

Specifically, heatmaps from the last conv-deconv stack together with their multi-scale heatmaps are concatenated and fed to the fully convolutional regression network. Thus the pose structure is refined by the regressing feature map across all feature scales and body keypoints. This regression process can effectively refine keypoint locations in considering body structural priors. Fig. 3(c,d) shows an example of our multi-scale, across-keypoint fine-tuning with improved keypoint heatmaps and pose estimation accuracy.

3. IMPLEMENTATION AND EXPERIMENTS

We train and test our model on two public datasets – MPII (28K/12k train/test) [17] and FLIC (5k/1k train/test) [5] respectively. Our stacked conv/deconv hourglass modules are

¹If we remove the multi-scale intermediate supervision (GT/2 and GT/4 in Fig. 2), and keep only the single-scale GT intermediate supervision (dark brown circle), as well as ignoring the regression network at the end, our network is reduced to an architecture similar to [11].




Fig. 3. Keypoint regression to disambiguate multiple peaks in the keypoint heatmaps. (a-b) shows an example of (a) the keypoint prediction and (b) the heatmaps from the conv-deconv module, which will be fed into the regression network. (c-d) shows (c) the output keypoint locations and (d) the heatmaps after regression. Observe that the heatmap peaks in (d) are more focused than in (b).

trained on the respective datasets using the ADAM optimizer for 100 epochs, starting with initial learning rate 0.0005 with decay. Evaluations are described in three subsections. §3.1 describes the accuracy evaluation on the two datasets. §3.2 reports experiments regarding our network design and parameters, including the number of hourglasses and multi-scale supervision to investigate their effects regarding performance. §3.3 evaluates how the multi-scale supervision can improve the handling of body part occlusions.

3.1. Evaluation on Accuracy

Evaluation is conducted using the standard Percentage of Correct Keypoints (PCK) metric [19], which reports the percentage of keypoint detection falling within a normalized distance of the ground truth. For FLIC, PCK is set to the percentage of disparities between the detected keypoints *w.r.t.* the groundtruth after a normalization against a fraction of the torso size. For MPII, such disparities are normalized by a fraction of the head size, which is denoted as PCK^h . The PCK evaluation metric is defined as:

$$PCK(\alpha) = \frac{1}{M} \frac{1}{N} \sum_{m=1}^M \sum_{n=1}^N I_A \left(\frac{\|Gt_n - Pred_n\|_2}{H} < \alpha \right), \quad (2)$$

where M is the dataset size, and N is the number of keypoints of a person. $I_A(\cdot)$ is an indicator function: I_A is 1 if (\cdot) is true, or 0 otherwise. $\|Gt_n - Pred_n\|_2$ is the Euclidean distance between the groundtruth and the prediction of the location of keypoint n . The normalization H is half of the head size for PCK^h and the torso size for PCK. Finally, α is the threshold to estimate if a keypoint is predicted correctly.


Table 1 summarizes the MPII performance evaluation. Observe that our method achieves state-of-the-art results across all keypoints (top 1 or 2, except the head) on the MPII dataset. Table 2 summarizes the FLIC results, where our PCK reaches 99.2% for the elbow, and 97.3% for the wrist. Our method performs better on shoulders, elbows, wrists that are in general harder to detect. This is due to improvements in our multi-scale feature supervision and global joint regression.

Table 1. Evaluation results on the MPII pose dataset ($PCK^h=0.5$)

	Head	Shoulder	Elbow	Wrist	Hip	Knee	Ankle	Total
Tompson <i>et al.</i> CVPR'15 [19]	96.1	91.9	83.9	77.8	80.9	72.3	64.8	82.0
Belagiannis & Zisserman FG'17 [20]	97.7	95.0	88.2	83.0	87.9	82.6	78.4	88.1
Insafutdinov <i>et al.</i> ECCV'16 [21]	96.8	95.2	89.3	84.4	88.4	83.4	78.0	88.5
Wei <i>et al.</i> CVPR'16 [11]	97.8	95.0	88.7	84.0	88.4	82.8	79.4	88.5
Bulat & Tzimiropoulos ECCV'16 [22]	97.9	95.1	89.9	85.3	89.4	85.7	81.7	89.7
Our model	97.0	95.8	90.9	86.3	89.1	85.0	80.8	89.8

Table 2. Results on the FLIC dataset ($PCK=0.2$)

	Elbow	Wrist
Tompson <i>et al.</i> CVPR'15 [9]	93.1	92.4
Wei <i>et al.</i> CVPR'16 [11]	97.8	95.0
Our model	99.2	97.3

**Fig. 4.** Performance comparisons on the number of multi-scale supervision and network depth. Observe that increasing network depth (number of conv-deconv stacks) results in significant performance boost, since deeper network can extract better features for keypoint detection. Performance also increases with the number of multi-scale supervisions.

3.2. Evaluation on Network Parameters

We evaluate the components of the multi-scale network in two aspects on the MPII validation set, as shown in Fig. 4: (1) the number of conv-deconv stacks used in the network, and (2) the number of scales with intermediate supervisions: (i) groundtruth scale-only (GT, i.e. the original intermediate supervision as in [11] (red line), (ii) GT and GT/2 (green line), (iii) GT, GT/2, and GT/4 (blue line).

For pose estimation, a deeper network can mostly outperform a shallow one. However the network depth is limited by the available computational resource, especially the GPU memory used during training. State-of-the-art works [11, 14, 15] use 4 GTX Titan X GPUs to run 8 conv-deconv stacks and 256 feature channels in the conv layers. In this paper, we use 4 conv-deconv stacks and 64 feature channels, in order to fit the model on a single GTX 1080 GPU. Our training resource is only 1/8 of the state-of-the-art works.

Fig. 4 shows that the increasing use of conv-deconv stacks can consistently improve performance, which is expectable. It also shows that the increasing use of number of multi-scale

**Fig. 5.** Pose detection results on selected challenging samples from the MPII test set. These scenarios contain cluttered background, heavy occlusions, and activities involving multiple people near the subject of interest.

supervision can consistently improve performance.

3.3. Evaluation on Occlusion Handling

Occlusion is a common challenge for human pose estimation. We evaluate our method on a subset of MPII test set with available occluded keypoint labels. We focus on occluded keypoints which are connected to and can be inferred from other visible body parts, e.g. a hidden elbow can be recovered from visible shoulder and wrist locations. This experiment can evaluate how the proposed structural regression network performs hand-in-hand with multi-scale feature supervision for occlusion recovery. We obtain 86.7% for $PCK^h=0.5$ with GT, GT/2, GT/4 multi-scale supervision. In comparison, the score is 84.3% without multi-scale supervision.

Fig. 5 show our results on a few challenge cases in MPII test set involving multiple persons and complex background/occlusions. Observe that the proposed method can produce plausible results.

4. CONCLUSION

We present an improved network with multi-scale supervision and structural keypoints regression for human pose estimation. We show that both improvements can consistently increase performance when comparing with state-of-the-art methods. Our method can effectively handle challenge cases including part occlusions, complex background and activities.

Future work includes the use of deeper network stacks on multiple GPUs aiming for multi-person scenarios.

5. REFERENCES

- [1] Zhao Liu, Jianke Zhu, Jiajun Bu, and Chun Chen, “A survey of human pose estimation,” *Journal of Visual Communication and Image Representation*, vol. 32, pp. 10–19, 2015.
- [2] Lubomir Bourdev and Jitendra Malik, “Poselets: Body part detectors trained using 3D human pose annotations,” in *ICCV*, 2009, pp. 1365–1372.
- [3] J. Charles, T. Pfister, D. Magee, D. Hogg, and A. Zisserman, “Domain adaptation for upper body pose tracking in signed TV broadcasts,” in *BMVC*, 2013, pp. 47.1–47.11.
- [4] A. Cherian, J. Mairal, K. Alahari, and C. Schmid, “Mixing body-part sequences for human pose estimation,” in *CVPR*, 2014, pp. 2361–2368.
- [5] Benjamin Sapp and Ben Taskar, “Multimodal decomposable models for human pose estimation,” in *CVPR*, 2013, pp. 3674–3681.
- [6] M.-C. Chang, H. Qi, X. Wang, H. Cheng, and S. Lyu, “Fast online upper body pose estimation from video,” in *BMVC*, Swansea, England, 2015, pp. 104.1–104.12.
- [7] Alexander Toshev and Christian Szegedy, “DeepPose: Human pose estimation via deep neural networks,” *CVPR*, pp. 1653–1660, 2014.
- [8] Tomas Pfister, James Charles, and Andrew Zisserman, “Flowing convnets for human pose estimation in videos,” in *ICCV*, 2015, pp. 1913–1921.
- [9] Jonathan J. Tompson, Arjun Jain, Yann LeCun, and Christoph Bregler, “Joint training of a convolutional network and a graphical model for human pose estimation,” in *NIPS* 27, 2014, pp. 1799–1807.
- [10] Xiao Chu, Wanli Ouyang, Hongsheng Li, and Xiaogang Wang, “Structured feature learning for pose estimation,” in *CVPR*, 2016, pp. 4715–4723.
- [11] Shiheng Wei, Varun Ramakrishna, Takeo Kanade, and Yaser Sheikh, “Convolutional pose machines,” *CVPR*, pp. 4724–4732, 2016.
- [12] Jonathan J. Tompson, Arjun Jain, Yann LeCun, and Christoph Bregler, “Joint training of a convolutional network and a graphical model for human pose estimation,” in *NIPS*, 2014, pp. 1799–1807.
- [13] Xiao Chu, Wanli Ouyang, Hongsheng Li, and Xiaogang Wang, “Structured feature learning for pose estimation,” in *CVPR*, 2016, pp. 4715–4723.
- [14] Alejandro Newell, Kaiyu Yang, and Jia Deng, “Stacked hourglass networks for human pose estimation,” in *ECCV*, 2016, pp. 483–499.
- [15] Xiao Chu, Wei Yang, Wanli Ouyang, Cheng Ma, Alan L. Yuille, and Xiaogang Wang, “Multi-context attention for human pose estimation,” in *CVPR*, 2017, pp. 5669–5678.
- [16] Wei Yang, Shuang Li, Wanli Ouyang, Hongsheng Li, and Xiaogang Wang, “Learning feature pyramids for human pose estimation,” in *ICCV*, 2017, pp. 1290–1299.
- [17] Mykhaylo Andriluka, Leonid Pishchulin, Peter Gehler, and Bernt Schiele, “2D human pose estimation: New benchmark and state of the art analysis,” in *CVPR*, 2014, pp. 3686–3693.
- [18] Zhe Cao, Tomas Simon, Shih-En Wei, and Yaser Sheikh, “Realtime multi-person 2D pose estimation using part affinity fields,” *arXiv preprint arXiv:1611.08050*, 2016.
- [19] Jonathan Tompson, Ross Goroshin, Arjun Jain, Yann LeCun, and Christopher Bregler, “Efficient object localization using convolutional networks,” *CVPR*, pp. 648–656, 2015.
- [20] Vasileios Belagiannis and Andrew Zisserman, “Recurrent human pose estimation,” *FG*, pp. 468–475, 2017.
- [21] Eldar Insafutdinov, Leonid Pishchulin, Bjoern Andres, Mykhaylo Andriluka, and Bernt Schiele, “DeeperCut: A deeper, stronger, and faster multi-person pose estimation model,” *ECCV*, pp. 34–50, 2016.
- [22] Adrian Bulat and Georgios Tzimiropoulos, “Human pose estimation via convolutional part heatmap regression,” *ECCV*, pp. 717–732, 2016.