

The Devil is in the Details: Delving into Unbiased Data Processing for Human Pose Estimation

Junjie Huang,¹ Zheng Zhu,² Feng Guo,¹ Guan Huang¹

¹XForwardAI Technology Co.,Ltd, Beijing, China

²Tsinghua University, Beijing, China

{junjie.huang, feng.guo, guan.huang}@xforwardai.com, zhengzhu@ieee.org

Abstract

Recently, the leading performance of human pose estimation is dominated by top-down methods. Being a fundamental component in training and inference, data processing has not been systematically considered in pose estimation community, to the best of our knowledge. In this paper, we focus on this problem and find that the devil of top-down pose estimator is in the biased data processing. Specifically, by investigating the standard data processing in state-of-the-art approaches mainly including data transformation and encoding-decoding, we find that the results obtained by common flipping strategy are unaligned with the original ones in inference. Moreover, there is statistical error in standard encoding-decoding during both training and inference. Two problems couple together and significantly degrade the pose estimation performance. Based on quantitative analyses, we then formulate a principled way to tackle this dilemma. Data is processed in continuous space based on unit length (the intervals between pixels) instead of in discrete space with pixel, and a combined classification and regression approach is adopted to perform encoding-decoding. The Unbiased Data Processing (UDP) for human pose estimation can be achieved by combining the two together. UDP not only boosts the performance of existing methods by a large margin but also plays a important role in result reproducing and future exploration. As a model-agnostic approach, UDP promotes SimpleBaseline-ResNet50-256 × 192 by 1.5 AP (70.2 to 71.7) and HRNet-W32-256 × 192 by 1.7 AP (73.5 to 75.2) on COCO test-dev set. The HRNet-W48-384 × 288 equipped with UDP achieves 76.5 AP and sets a new state-of-the-art for human pose estimation. The source code is publicly available for further research¹.

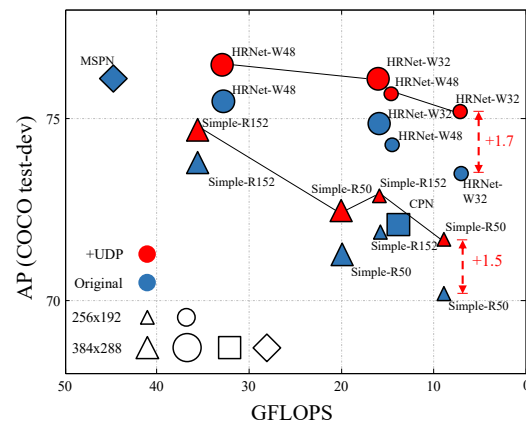


Figure 1. The improvement of AP on COCO test-dev set when the proposed Unbiased Data Processing (UDP) is applied.

1. Introduction

Human pose estimation is of importance for visual understanding tasks such as video surveillance [15] and action recognition [4, 37, 36]. In recent years, research community has witnessed a significant advance from single person [2, 10, 30, 29, 31, 21, 33] to multi-person pose estimation [24, 13, 3, 23, 6, 26, 20, 7], where the latter can be generally categorized into bottom-up [24, 13, 3, 20, 22, 7] and top-down approaches [23, 6, 11, 32, 26].

While most state-of-the-art top-down methods [6, 32, 26, 16] focus on the design of the network structure, we pay attention to the *data processing* aspect considering it as another fundamental component. All visual recognition tasks are born with data processing, and in general share data processing methodology with each other like data augmentation and transformation between different coordinate systems. However, when compared with other tasks like classification [25], object detection [17] and semantic segmentation [19, 8], the performance of human pose estimation algorithms is much more sensitive to the methods used in data processing on account of the evaluation principle. In

¹<https://github.com/HuangJunJie2017/UDP-Pose>

the evaluation of human pose estimation, the metrics are calculated based on the positional offset between ground truth labels and prediction results [17, 1], where small systematic bias caused by data processing will degrade the performance of pose estimators.

To the best of our knowledge, data processing has not been systematically considered in human pose estimation community. When this topic is addressed, we find that most state-of-the-art systems [6, 32, 16, 26] suffer from the same two common problems: i) the unaligned results in inference obtained by using flipping strategy, which is derived from analyzing this problems in discrete space and utilizing *pixel* to measure the size of images in data transformation. ii) statistical error in standard encoding-decoding during the training and inference, respectively.

We take the biased data processing used in SimpleBaseline [32] and HRNet [26] as an example to address these problems. Quantitatively, we show that the aforementioned two problems couple together and significantly degrade the pose estimation performance. Based on the analysis results, a principled Unbiased Data Processing (UDP) is proposed to tackle this dilemma. Specifically, we analyze the data transformation in continuous space and use unit length (the intervals between pixels) instead of pixel for image size measurement, leading to aligned pose results when flipping is performed in inference. Moreover, motivated by the combined classification and regression pose estimator [23], a theoretically error-free encoding-decoding method cooperating with our data transformation is designed to further promote the performance. It is worth noting that UDP is a model-agnostic approach, which can serve for most top-down pipelines. On challenging COCO human pose estimation dataset, the proposed UDP sets new state-of-the-art performance among competitors. As shown in Figure 1, on COCO *test-dev* set, our approach promotes SimpleBaseline by 1.5 AP (70.2 to 71.7) and 1.0 AP (71.9 to 72.9) within ResNet50-256 × 192 and ResNet152-256 × 192 configurations, respectively. For HRNet within *W*32-256 × 192 and *W*48-256 × 192 configurations, UDP obtains gains by 1.7 AP (73.5 to 75.2) and 1.4 AP (74.3 to 75.7), respectively. The HRNet-*W*48-384 × 288 equipped with UDP achieves 76.5 AP (1.0 improvement) and sets a new state-of-the-art for human pose estimation.

The main contributions of this paper can be summarized as follows:

1. This paper quantitatively analyzes the common biased data processing for human pose estimation. Interestingly, we find that the systematic error in standard data transformation and encoding-decoding couple together and significantly degrade the performance of top-down pipelines. To the best of our knowledge, this is the first work to systematically address the data processing in pose community.

2. Based on analysis results, this paper formulates a principled Unbiased Data Processing (UDP) strategy, which equips with *unit length*-based measurement and combined classification and regression encoding-decoding. The proposed UDP is a model-agnostic strategy and can be utilized in most top-down pose estimators. We hope UDP will be significant for result reproducing and future research.
3. On challenging COCO human pose estimation dataset, UDP promotes state-of-the-arts by large margin among variable backbones and input sizes. Specifically, the HRNet-*W*48-384 × 288 equipped with UDP achieves 76.5 AP on COCO *test-dev* set and sets a new state-of-the-art for human pose estimation. It is worth noting that our approach only increases negligible calculation burden during training and inference.

2. Related Work

Bottom-up methods start by detecting identity-free joints for all the persons in an input image through predicting heatmaps of different classes keypoints and then group them into person instances. OpenPose [3] builds a model that contains two branches to predict keypoint heatmaps and pairwise relationships (part affinity fields) between them. Newell et al. [20] use one network for both heatmap prediction and grouping. Grouping is done by association embedding, which assigns each keypoint with a tag and groups keypoints based on the L2 distance between tag vectors. MultiPoseNet [14] simultaneously achieves human detection and pose estimation, and proposes PRN to group the keypoints by the bounding box of each people. HigherHRNet [7] maintains high-resolution feature maps which effectively improves the precision of the predictions.

Top-down methods achieve multi-person pose estimation by the two-stages process, including obtaining person bounding boxes by a person detector and predicting keypoint locations within these boxes. CPN [6] and MSPN [16] are the leading methods on COCO keypoint challenge, adopting cascade network to refine the keypoints prediction. SimpleBaseline [32] adds a few deconvolutional layers to enlarge the resolution of output features. It is simple but effective in performance improvement. HRNet [26] maintains high-resolution representations through the whole process, achieving state-of-the-art performance on public dataset. Mask R-CNN [11] builds an end-to-end framework and achieves a good balance between performance and inference speed. As single person pose estimation is performed with fixed scale patches, most state-of-the-art performances on multi-person popular benchmarks are achieved by top-down methods.

Data processing in human pose estimation with top-down paradigm mainly includes data transformation, data

augmentation and encoding-decoding. **Data transformation** means transforming the keypoint location between different coordinate systems such as source image, network input and output. During this process, most state-of-the-art methods [6, 32, 16, 26] use *pixel* to measure the size of images, leading to unaligned results when using flipping strategy in inference. [32, 26] empirically shift the result from flipped image by 1 pixel in network output coordinate system to suppress the predicting error. [6, 16] achieve similar effect by shifting the average result by 2 pixels in network input coordinate system. These compensations are effective but limited. **Data augmentation** is a common strategy for increasing the diversity of samples, which can do help to enhance the robustness of the algorithms. Common augmentation strategies include random rotation, random scale, flip and half body [26]. All data augmentations are performed in the data transformation from the source image into the network input. Finally, **encoding-decoding** refers to transforming between joint coordinates and heatmaps, which is firstly proposed in [29] and has been widely used in state-of-the-art methods [11, 7, 6, 32, 16, 26]. In training process, they encode the ground truth into a heatmap with Gaussian distribution centered at the keypoint position. And decoding means transforming the network predicted heatmap back into keypoint coordinate in inference process. This pipeline shows superior performance when compared with directly predicting the keypoint coordinates [27], but introduces a systematic error that degrades the accuracy of predictions. In contrast, the combined classification and regression encoding-decoding paradigm [23] provides an error-free entrance to further promote the prediction accuracy of top-down methods.

3. Unbiased Data Processing for Human Pose Estimation

In this section, we analyze the standard data processing approaches in current state-of-the-arts from two aspects: data transformation and encoding-decoding. Following the analysis, the unbiased data processing strategy is introduced to effectively promote the performance of pose estimators.

Symbol Definition In this paper, three coordinate systems are adopted: source image coordinate systems (*i.e.* original image coordinate systems, denoted as $O_s-X_s-Y_s$), network input coordinate systems (*i.e.* cropped/resized image coordinate systems, denoted as $O_i-X_i-Y_i$), and network output coordinate systems (*i.e.* heatmap coordinate systems, denoted as $O_o-X_o-Y_o$). These three coordinate systems define corresponding three spaces: source image space (denoted by superscript s), network input image space (denoted by superscript i), and network output image space (denoted by superscript o). The heatmaps related to encoding-decoding

are defined in network output image space. In following parts, we use superscript p to denote that the length is measured in *pixel*. Otherwise the length is measured in *unit length* related to the corresponding space. The image matrix and corresponding keypoint coordinates are denoted by \mathbf{I} and \mathbf{k} , respectively. The symbol with hat $\hat{\mathbf{k}}$ is the network predicted results of the corresponding ground truth label \mathbf{k} .

3.1. Data Transformation

3.1.1 Analysis of the Standard Data Transformation

The data transformation means transforming the keypoint location such as cropping, rotation, resizing and flipping between different coordinate systems. Existing pose estimation methods adopt *pixel* to measure the size of the images, which is in a discrete space. However, for positioning task, *pixels* are some sample points in the image plane (*i.e.* continuous space). For example, if the size of an image is $({}^p w, {}^p h)$ measured by *pixel*, its size in the continuous image plane is $({}^p w - 1, {}^p h - 1)$. Using the *pixel* as the measurement would significantly degrade the performance when the *de facto* standard flipping strategy is performed during inference [32, 26].

During the training process, top-down pipelines firstly transform the source image sample into an augmented sample in the network input space. We denote the network input image matrix as ${}^i I$ with size of $({}^p w_i, {}^p h_i)$ and the corresponding keypoint coordinates as ${}^i \mathbf{k}$. As illustrated in Figure 2, when mapping network input to the source image space, the size of the contents in the source image can be formulated into a bounding box (The green boxes in source image space in Figure 2) in $O_s-X_s-Y_s$ coordinate system with center $({}^s x_b, {}^s y_b)$ and scale $({}^s w_b, {}^s h_b)$. Then ${}^i \mathbf{k}$ obtained by the standard methods [32, 26] can be formulated as:

$${}^i \mathbf{k} = \begin{bmatrix} \frac{{}^p w_i c\theta}{{}^s w_b} & -\frac{{}^p w_i s\theta}{{}^s w_b} & \frac{{}^p w_i (-{}^s x_b c\theta + {}^s y_b s\theta + 0.5{}^s w_b)}{{}^s w_b} \\ \frac{{}^p h_i s\theta}{{}^s h_b} & \frac{{}^p h_i c\theta}{{}^s h_b} & \frac{{}^p h_i (-{}^s x_b s\theta - {}^s y_b c\theta + 0.5{}^s h_b)}{{}^s h_b} \\ 0 & 0 & 1 \end{bmatrix} {}^s \mathbf{k} \quad (1)$$

where $c\theta$ and $s\theta$ denote $\cos(\theta)$ and $\sin(\theta)$, respectively. θ is the angle in rotation augmentation. The details inference of this transformation can be found in Appendix. And each pixel in network input matrix can be backtracked to the source image as follows:

$${}^s \mathbf{x} = \begin{bmatrix} \frac{{}^s w_b c\theta}{{}^p w_i} & \frac{{}^s h_b s\theta}{{}^p h_i} & -0.5{}^s w_b c\theta - 0.5{}^s h_b s\theta + {}^s x_b \\ -\frac{{}^s w_b s\theta}{{}^p w_i} & \frac{{}^s h_b c\theta}{{}^p h_i} & 0.5{}^s w_b s\theta - 0.5{}^s h_b c\theta + {}^s y_b \\ 0 & 0 & 1 \end{bmatrix} {}^i \mathbf{x} \quad (2)$$

where ${}^i \mathbf{x}$ is the coordinate of each pixel in image matrix ${}^i I$. As illustrated in Figure 2, the standard augmented sample can be equivalently regarded as the result produced by the following two steps:

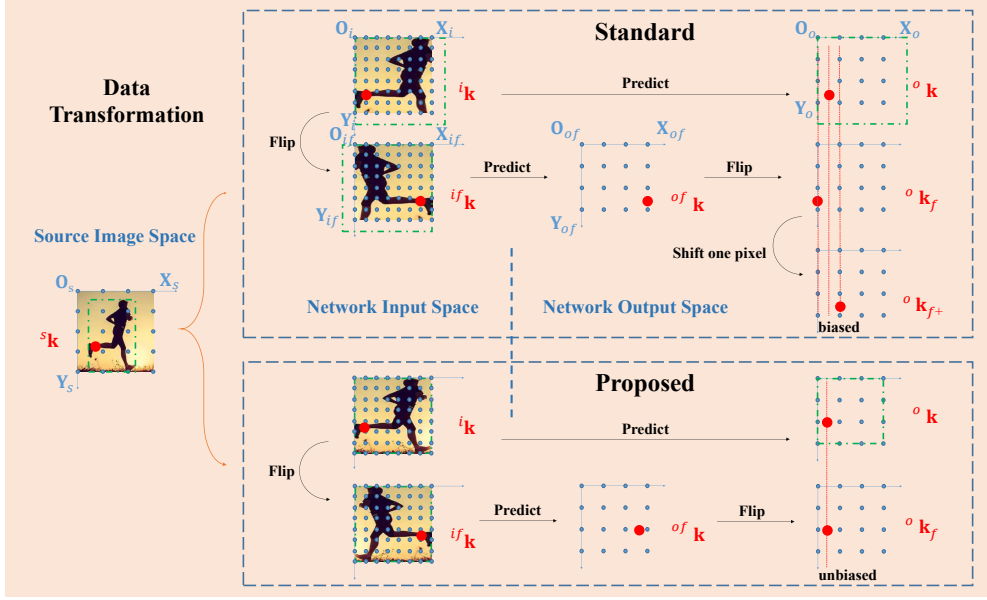


Figure 2. Illustration for the processes of standard biased data transformation and proposed unbiased data transformation. The network input size is assumed to be $(^p w_i, ^p h_i) = (8, 8)$ and the stride factor s is assumed to be 2.

1. Cropping the region of interest from the source image and resizing it into the shape of $(^p w_i + 1, ^p h_i + 1)$ (The green boxes in network input space in Figure 2).
2. Cropping the aforementioned result by 1 pixel both on the right edge and the bottom edge, resulting in an image with the size of $(^p w_i, ^p h_i)$.

The training sample generated by this method is semantically aligned with the original sample (*i.e.* the pose annotations are still in the proper positions). When producing the ground truth in heatmap whose size is $(^p w_o, ^p h_o)$, the standard methods transform the input keypoint positions by the stride factor $s = ^p w_i / ^p w_o = ^p h_i / ^p h_o$:

$${}^o \mathbf{k} = \frac{1}{s} {}^i \mathbf{k} \quad (3)$$

where ${}^o \mathbf{k}$ is the keypoint coordinates in network output heatmap. During the training process, the network learns the pattern that infers a response map centered at ${}^o \mathbf{k}$ position according to the input image contents.

In inference stage, the standard method maps the predicted result ${}^o \hat{\mathbf{k}}$ to the source image space by:

$${}^s \hat{\mathbf{k}} = \begin{bmatrix} \frac{{}^s w_b c \theta}{{}^p w_o} & \frac{{}^s h_b s \theta}{{}^p h_o} & -0.5 {}^s w_b c \theta - 0.5 {}^s h_b s \theta + {}^s x_b \\ -\frac{{}^s w_b s \theta}{{}^p w_o} & \frac{{}^s h_b c \theta}{{}^p h_o} & 0.5 {}^s w_b s \theta - 0.5 {}^s h_b c \theta + {}^s y_b \\ 0 & 0 & 1 \end{bmatrix} {}^o \hat{\mathbf{k}} \quad (4)$$

where ${}^s \hat{\mathbf{k}}$ is the final predicted position of keypoints in the source image space. This transformation also can be equivalently formulated into two steps:

1. Padding the size of network output heatmap from $(^p w_o, ^p h_o)$ into $(^p w_o + 1, ^p h_o + 1)$ (The green boxes in network output space in Figure 2),
2. Mapping the aforementioned padded heatmap to the corresponding bounding box in the source image.

Ideally, the predicted result ${}^s \hat{\mathbf{k}}$ is equal to ${}^s \mathbf{k}$. However, the result would be biased when the flipping strategy is utilized. Standard methods flip the network input images, then the keypoint ${}^i \mathbf{k}$ is located at:

$${}^i f \mathbf{k} = \begin{bmatrix} -1 & 0 & ^p w_i - 1 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} {}^i \mathbf{k} \quad (5)$$

where ${}^i f \mathbf{k}$ is the corresponding location of keypoint ${}^i \mathbf{k}$ in flipped image. According to Equation 3, standard methods predict keypoint in the output heatmap as:

$${}^o f \hat{\mathbf{k}} = \frac{1}{s} {}^i f \mathbf{k} \quad (6)$$

Then, the final result of flipped image ${}^o \hat{\mathbf{k}}_f$ can be obtained by flipping back:

$${}^o \hat{\mathbf{k}}_f = \begin{bmatrix} -1 & 0 & ^p w_o - 1 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} {}^o f \hat{\mathbf{k}} = \begin{bmatrix} 1 & 0 & -\frac{s-1}{s} \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} {}^o \hat{\mathbf{k}} \quad (7)$$

Here, ${}^o \hat{\mathbf{k}}_f$ is not exactly aligned with ${}^o \hat{\mathbf{k}}$, and there is an offset of $-\frac{s-1}{s}$ in O_o - X_o direction. If we directly average

${}^o\hat{\mathbf{k}}$ and ${}^o\hat{\mathbf{k}}_f$ as reported in [32, 26]:

$${}^o\hat{\mathbf{k}}_a = \frac{{}^o\hat{\mathbf{k}}_f + {}^o\hat{\mathbf{k}}}{2} \quad (8)$$

the corresponding error in O_o - X_o direction is:

$${}^oe(x) = |x({}^o\mathbf{k}) - x({}^o\hat{\mathbf{k}}_a)| = \left| -\frac{s-1}{2s} \right| \quad (9)$$

Standard approaches [32, 26] shift the flipped result by 1 pixel before averaging operation to narrow this gap:

$${}^o\hat{\mathbf{k}}_{f+} = \begin{bmatrix} 1 & 0 & 1 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} {}^o\hat{\mathbf{k}}_f = \begin{bmatrix} 1 & 0 & \frac{1}{s} \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} {}^o\hat{\mathbf{k}} \quad (10)$$

In this way, the final error can be reduced to ${}^oe(x)' = \left| \frac{1}{2s} \right|$. ${}^oe(x)' < {}^oe(x)$ when $s > 2$, which makes sense in most existing methods. Intuitively, a compensation for ${}^oe(x)'$ can make the result more accurate. However, the actual contribution of this direct compensation is very limited. We will give detailed analysis of this inefficiency in Section 3.2. Besides, when mapping ${}^oe(x)'$ back to source image coordinate system (O_s - X_s - Y_s) and considering Equation 4 with $\theta = 0$, we have:

$${}^se(x)' = \left| \frac{1}{2s} \times \frac{{}^sw_b}{p w_i} \right| = \left| \frac{{}^sw_b}{2^p w_i} \right| \quad (11)$$

where sw_b is fixed in inference process. So larger network input size can help suppress the predicted error caused by ${}^oe(x)'$. In other words, the standard methods benefit more from higher input resolution and suffer more accuracy loss from lower input resolution.

3.1.2 The Proposed Data Transformation

In this paper, a principled way is proposed to address the misalignment problem. Specifically, we adopt *unit length* as the image size measurement criterion, which is defined as the distance between two adjacent pixels in a specific space. Based on this concept, the ground truth label ${}^i\mathbf{k}$ in network input space should be obtained by the following transformation:

$${}^i\mathbf{k} = \begin{bmatrix} \frac{p w_i - 1}{s w_b / c\theta} & -\frac{p w_i - 1}{s w_b / s\theta} & \frac{p w_i - 1}{s w_b} (-s x_b c\theta + s y_b s\theta + 0.5 s w_b) \\ \frac{p h_i - 1}{s h_b / s\theta} & \frac{p h_i - 1}{s h_b / c\theta} & \frac{p h_i - 1}{s h_b} (-s x_b s\theta - s y_b c\theta + 0.5 s h_b) \\ 0 & 0 & 1 \end{bmatrix} {}^s\mathbf{k} \quad (12)$$

And each pixel in network input image matrix should be backtracked to the source image space by:

$${}^s\mathbf{x} = \begin{bmatrix} \frac{s w_b c\theta}{p w_i - 1} & \frac{s h_b s\theta}{p h_i - 1} & -0.5 s w_b c\theta - 0.5 s h_b s\theta + s x_b \\ -\frac{s w_b s\theta}{p w_i - 1} & \frac{s h_b c\theta}{p h_i - 1} & 0.5 s w_b s\theta - 0.5 s h_b c\theta + s y_b \\ 0 & 0 & 1 \end{bmatrix} {}^i\mathbf{x} \quad (13)$$

While producing the ground truth label on heatmap, we should measure the map in continuous space and utilize factor $t = (p w_i - 1)/(p w_o - 1) = (p h_i - 1)/(p h_o - 1)$:

$${}^o\mathbf{k} = \frac{1}{t} {}^i\mathbf{k} \quad (14)$$

In this way, the result of flipped image ${}^o\hat{\mathbf{k}}_f$ is exactly aligned with the original result ${}^o\hat{\mathbf{k}}$. Finally the prediction ${}^s\hat{\mathbf{k}}$ in source image space should be obtained by the following inverse transformation:

$${}^s\hat{\mathbf{k}} = \begin{bmatrix} \frac{s w_b c\theta}{p w_o - 1} & \frac{s h_b s\theta}{p h_o - 1} & -0.5 s w_b c\theta - 0.5 s h_b s\theta + s x_b \\ -\frac{s w_b s\theta}{p w_o - 1} & \frac{s h_b c\theta}{p h_o - 1} & 0.5 s w_b s\theta - 0.5 s h_b c\theta + s y_b \\ 0 & 0 & 1 \end{bmatrix} {}^o\hat{\mathbf{k}} \quad (15)$$

3.2. Encoding and Decoding

The aforementioned analysis is given under the precondition that the encoding-decoding process between keypoint position and heatmap is precise (*i.e.* $\hat{\mathbf{k}} = \mathbf{k}$). However this precondition doesn't stand in the standard methods [32, 26]. In the following we will firstly study the systematic error (*i.e.* $|\hat{\mathbf{k}} - \mathbf{k}|$) in the standard encoding-decoding and show how this systematic error affects the aforementioned conclusions. Only the network output heatmap coordinate system O_o - X_o - Y_o is used in this subsection.

3.2.1 Analysis of the Standard Encoding-decoding

The standard encoding method. Formally, given the ground truth label point $\mathbf{k} = (m, n)$ in the heatmap, [26, 32] firstly quantize the label point coordinates to obtain integer type label coordinates \mathbf{k}_q :

$$\mathbf{k}_q = (m_q, n_q) = \mathcal{R}(\mathbf{k}) = (\mathcal{R}(m), \mathcal{R}(n)) \quad (16)$$

where \mathcal{R} denotes the rounding operation. Then, the heatmap centering at \mathbf{k}_q is generated by:

$$\mathcal{H}(x, y, \mathbf{k}_q) = \exp\left(-\frac{(x - m_q)^2 + (y - n_q)^2}{2\delta^2}\right) \quad (17)$$

where (x, y) denotes the coordinate of each element in the heatmap, and δ denotes a fixed spatial variance.

The standard decoding method. Given a trained network prediction $\hat{\mathcal{H}}(x, y, \hat{\mathbf{k}})$ and under the ideal condition of $\hat{\mathcal{H}} = \mathcal{H}$, [26, 32] decode it by firstly locating the highest response:

$$\hat{\mathbf{k}}_q = (\hat{m}_q, \hat{n}_q) = \operatorname{argmax}(\hat{\mathcal{H}}) \quad (18)$$

According to Equation 17 and Equation 18, the predicted keypoint position in the x direction of heatmap is:

$$\hat{m}_q = \begin{cases} \mathcal{F}(m) & \text{if } m - \mathcal{F}(m) < 0.5 \\ \mathcal{C}(m) & \text{otherwise} \end{cases} \quad (19)$$

where \mathcal{F} and \mathcal{C} denote the floor operation and ceil operation, respectively. As the ideal location of $\hat{\mathbf{k}}_q$ is \mathbf{k}_q , with the assumption that \mathbf{k} is uniformly distributed in the image plane, the expected error in each direction is $E(|m - \hat{m}_q|) = E(|n - \hat{n}_q|) = 1/4$ unit length with variance of $V(|m - \hat{m}_q|) = V(|n - \hat{n}_q|) = 1/48$. To reduce this error, [26, 32] shift $\hat{\mathbf{k}}_q$ by 0.25 unit length in each direction according to the gradient of response:

$$\hat{\mathbf{k}} = \hat{\mathbf{k}}_q + 0.25\nabla(\hat{\mathcal{H}})|_{\mathbf{x}=\hat{\mathbf{k}}_q} \quad (20)$$

where ∇ is the gradient operator. This operation changes the decoding result distribution into:

$$\hat{m} = \begin{cases} \mathcal{F}(m) + 0.25 & \text{if } m - \mathcal{F}(m) < 0.5 \\ \mathcal{C}(m) - 0.25 & \text{otherwise} \end{cases} \quad (21)$$

and the expected error in each direction can be reduced to $E(|m - \hat{m}|) = E(|n - \hat{n}|) = 1/8$ unit length with variance of $V(|m - \hat{m}|) = V(|n - \hat{n}|) = 1/192 \approx 0.0052$.

When mapping $E(|m - \hat{m}|)$ back to source image coordinate system (O_s - X_s Y_s) and considering Equation 4 with $\theta = 0$, we have:

$$E(|^s m - ^s \hat{m}|) = E(|^o m - ^o \hat{m}|) \times \frac{{}^s w_b s}{{}^p w_i} \quad (22)$$

Considering error $E(|m - \hat{m}|)$ and $E(|n - \hat{n}|)$, the standard methods [32, 26] with fixed stride factor s can also benefit from higher network input resolution.

Test error ${}^o e(x)' = \frac{1}{2s}$ caused by inference flipping in Section 3.1 has impact on the error distribution. With a specific stride factor $s = 4$ [32, 26], the predicted heatmap is changed into $\hat{\mathcal{H}}' = \mathcal{H}(x, y, m + 0.125, n + 0.125)$, leading to the decoding result distribution as:

$$\hat{m}' = \begin{cases} \mathcal{F}(m) + 0.25 & \text{if } m - \mathcal{F}(m) < 0.375 \\ \mathcal{C}(m) - 0.25 & \text{if } 0.375 \leq m - \mathcal{F}(m) < 0.875 \\ \mathcal{C}(m) + 0.25 & \text{otherwise} \end{cases} \quad (23)$$

and the expected error in O_o - X_o direction is enlarged by just $1/32$ unit length to $E(|m - \hat{m}'|) = 5/32$ with a variance of $V(|m - \hat{m}'|) = 37/3072 \approx 0.012$. The decoding error is dominant and the data transformation error has little impact on the final predicting performance. If we only fixed the data transformation error, probably the contribution is less than the destabilization in training process.

From a statistical point of view, the aforementioned direct compensation on \hat{m}' by $-\frac{1}{2s}$ unit length can completely eliminate the effect caused by ${}^o e(x)'$, as $E(|m - (\hat{m}' - 0.125)|) = 1/8$ exactly equals to $E(|m - \hat{m}|)$, and $V(|m - (\hat{m}' - 0.125)|) = 1/192$ exactly equals to $V(|m - \hat{m}|)$. This compensation is effective but limited, we still need other encoding-decoding method to eliminate the statistical error caused by the standard approaches.

3.2.2 The Proposed Encoding-decoding

Inspired by [23], we introduce a combined classification and regression encoding-decoding method, whose expected value of error is zero. Each ground truth label point $\mathbf{k} = (m, n)$ is encoded into one heatmap:

$$\mathcal{H}(x, y, \mathbf{k}) = \begin{cases} 1 & \text{if } (x - m)^2 + (y - n)^2 < R \\ 0 & \text{otherwise} \end{cases} \quad (24)$$

and two offset maps:

$$\mathcal{X}(x, y, \mathbf{k}) = m - x \quad (25)$$

$$\mathcal{Y}(x, y, \mathbf{k}) = n - y \quad (26)$$

During decoding, we firstly use a Gaussian kernel K to filter the heatmap to make the highest response located around the ground truth label point. Then the highest score position is located by:

$$\hat{\mathbf{k}}_h = \operatorname{argmax}(\hat{\mathcal{H}} \otimes K) \quad (27)$$

where \otimes denotes the convolutional operation and K is a kernel generated according to the radial based function:

$$K(x, y, N, \sigma) = \frac{\exp(-\frac{(x-N)^2 + (y-N)^2}{2\sigma^2})}{\sum_x \sum_y \exp(-\frac{(x-N)^2 + (y-N)^2}{2\sigma^2})} \quad (28)$$

where the kernel size is $2N + 1$. Finally the coordinate is modified according to the offset maps:

$$\hat{\mathbf{k}} = \hat{\mathbf{k}}_h + [\hat{\mathcal{X}}(\hat{\mathbf{k}}_h, \mathbf{k}) \otimes K, \hat{\mathcal{Y}}(\hat{\mathbf{k}}_h, \mathbf{k}) \otimes K]^T \quad (29)$$

where K is used for smoothing the two predicted offset maps. In this way, the predicted result $\hat{\mathbf{k}}$ is theoretically equal to the ground-truth \mathbf{k} with an expected error of zero in ideal conditions (i.e. $\hat{\mathcal{H}} = \mathcal{H}$ and $\hat{\mathcal{X}} = \mathcal{X}$ and $\hat{\mathcal{Y}} = \mathcal{Y}$).

4. Experiments

4.1. Implementation Details

Our model is trained on COCO train subset, which is equipped about 57,000 images and 150,000 person instances. We evaluate the proposed UDP on the val set and test-dev set, containing about 5,000 images and 20,000 images, respectively. The AP evaluation metric is reported based on Object Keypoint Similarity (OKS). We set the training configuration strictly following [32, 26] for comparison. During inference, HTC [5] detector is used to detect human instances. With multi-scale test, the 80-class and person AP on COCO val set [17] are 52.9 and 65.1, respectively. The results of HRNet [26] and SimpleBaseline [32] on COCO val set with this human detection are reproduced for fair comparison. We report the performance of single model, and only flipping test strategy is used.

Method	Backbone	Pretrain	Input size	#Params	GFLOPS	AP	AP ⁵⁰	AP ⁷⁵	AP ^M	AP ^L	AR
Hourglass[21]	8-stage Hourglass	N	256 × 192	25.1M	14.3	66.9	-	-	-	-	-
CPN[6]	ResNet-50	Y	256 × 192	27.0M	6.20	69.4	-	-	-	-	-
CPN[6]	ResNet-50	Y	384 × 288	102M	13.9	71.6	-	-	-	-	-
MSPN[16]	MSPN	Y	256 × 192	120M	19.9	75.9	-	-	-	-	-
SimpleBaseline[32]	ResNet-50	Y	256 × 192	34.0M	8.90	71.3	89.9	78.9	68.3	77.4	76.9
+UDP	ResNet-50	Y	256 × 192	34.2M	8.96	72.9(+1.6)	90.0	80.2	69.7	79.3	78.2
SimpleBaseline[32]	ResNet-152	Y	256 × 192	68.6M	15.7	72.9	90.6	80.8	69.9	79.0	78.3
+UDP	ResNet-152	Y	256 × 192	68.8M	15.8	74.3(+1.4)	90.9	81.6	71.2	80.6	79.6
SimpleBaseline[32]	ResNet-50	Y	384 × 288	34.0M	20.0	73.2	90.7	79.9	69.4	80.1	78.2
+UDP	ResNet-50	Y	384 × 288	34.2M	20.1	74.0(+0.8)	90.3	80.0	70.2	81.0	79.0
SimpleBaseline[32]	ResNet-152	Y	384 × 288	68.6M	35.6	75.3	91.0	82.3	71.9	82.0	80.4
+UDP	ResNet-152	Y	384 × 288	68.8M	35.7	76.2(+0.9)	90.8	83.0	72.8	82.9	81.2
HRNet[26]	HRNet-W32	Y	256 × 192	28.5M	7.10	75.6	91.9	83.0	72.2	81.6	80.5
+UDP	HRNet-W32	Y	256 × 192	28.7M	7.16	76.8(+1.2)	91.9	83.7	73.1	83.3	81.6
HRNet[26]	HRNet-W48	Y	256 × 192	63.6M	14.6	75.9	91.9	83.5	72.6	82.1	80.9
+UDP	HRNet-W48	Y	256 × 192	63.8M	14.7	77.2(+1.3)	91.8	83.7	73.8	83.7	82.0
HRNet[26]	HRNet-W32	Y	384 × 288	28.5M	16.0	76.7	91.9	83.6	73.2	83.2	81.6
+UDP	HRNet-W32	Y	384 × 288	28.7M	16.1	77.8(+1.1)	91.7	84.5	74.2	84.3	82.4
HRNet[26]	HRNet-W48	Y	384 × 288	63.8M	32.9	77.1	91.8	83.8	73.5	83.5	81.8
+UDP	HRNet-W48	Y	384 × 288	63.6M	33.0	77.8(+0.7)	92.0	84.3	74.2	84.5	82.5

Table 1. Comparisons on COCO val set. #Params and FLOPS are calculated only for the pose estimation network. UDP denotes the proposed Unbiased Data Processing methods.

Method	Backbone	Input size	#Params	GFLOPS	AP	AP ⁵⁰	AP ⁷⁵	AP ^M	AP ^L	AR
Mask-RCNN[11]	ResNet-50-FPN	-	-	-	63.1	87.3	68.7	57.8	71.4	-
Integral Pose Regression[28]	ResNet-101	256 × 256	45.0M	11.0	67.8	88.2	74.8	63.9	74.0	-
G-RMI+extra data[23]	ResNet-101	353 × 257	42.6M	57.0	68.5	87.1	75.5	65.8	73.3	73.3
CPN[6]	ResNet-Inception	384 × 288	-	-	72.1	91.4	80.0	68.7	77.2	78.5
RMPE[9]	PyraNet[34]	320 × 256	28.1M	26.7	72.3	89.2	79.1	68.0	78.6	-
CFN[12]	-	-	-	-	72.6	86.1	69.7	78.3	64.1	-
CPN(ensemble)[6]	ResNet-Inception	384 × 288	-	-	73.0	91.7	80.9	69.5	78.1	79.0
Posefix[18]	ResNet-152	384 × 288	68.6M	35.6	73.6	90.8	81.0	70.3	79.8	79.0
CSANet[35]	ResNet-152	384 × 288	-	-	74.5	91.7	82.1	71.2	80.2	80.7
MSPN[16]	MSPN	384 × 288	120M	44.8	76.1	93.4	83.8	72.3	81.5	81.6
SimpleBaseline[32]	ResNet-50	256 × 192	34.0M	8.90	70.2	90.9	78.3	67.1	75.9	75.8
+UDP	ResNet-50	256 × 192	34.2M	8.96	71.7(+1.5)	91.1	79.6	68.6	77.5	77.2
SimpleBaseline[32]	ResNet-50	384 × 288	34.0M	20.0	71.3	91.0	78.5	67.3	77.9	76.6
+UDP	ResNet-50	384 × 288	34.2M	20.1	72.5(+1.2)	91.1	79.7	68.8	79.1	77.9
SimpleBaseline[32]	ResNet-152	256 × 192	68.6M	15.8	71.9	91.4	80.1	68.9	77.4	77.5
+UDP	ResNet-152	256 × 192	68.8M	15.9	72.9(+1.0)	91.6	80.9	70.0	78.5	78.4
SimpleBaseline[32]	ResNet-152	384 × 288	68.6M	35.6	73.8	91.7	81.2	70.3	80.0	79.1
+UDP	ResNet-152	384 × 288	68.8M	35.7	74.7(+0.9)	91.8	82.1	71.5	80.8	80.0
HRNet[26]	HRNet-W32	256 × 192	28.5M	7.10	73.5	92.2	82.0	70.4	79.0	79.0
+UDP	HRNet-W32	256 × 192	28.7M	7.16	75.2(+1.7)	92.4	82.9	72.0	80.8	80.4
HRNet[26]	HRNet-W32	384 × 288	28.5M	16.0	74.9	92.5	82.8	71.3	80.9	80.1
+UDP	HRNet-W32	384 × 288	28.7M	16.1	76.1(+1.2)	92.5	83.5	72.8	82.0	81.3
HRNet[26]	HRNet-W48	256 × 192	63.6M	14.6	74.3	92.4	82.6	71.2	79.6	79.7
+UDP	HRNet-W48	256 × 192	63.8M	14.7	75.7(+1.4)	92.4	83.3	72.5	81.4	80.9
HRNet[26]	HRNet-W48	384 × 288	63.6M	32.9	75.5	92.5	83.3	71.9	81.5	80.5
+UDP	HRNet-W48	384 × 288	63.8M	33.0	76.5(+1.0)	92.7	84.0	73.0	82.4	81.6

Table 2. The improvement of AP on COCO test-dev set when the proposed UDP is applied to state-of-the-art methods.

4.2. Comparison with State-of-the-arts

Results on the val set. The results of proposed method and state-of-the-arts are listed in Table 1. We report the performance improvement when UDP is applied to the recent state-of-the-art human pose estimation methods SimpleBaseline [32] and HRNet [26]. For SimpleBaseline with ResNet-50, the promotion is 1.6 AP (71.3 to 72.9) and 1.4

AP (72.9 to 74.3). For larger backbone, the promotion is 0.8 and 0.9 respectively. For HRNet with different backbones and input sizes, improvements of the proposed UDP are 1.2 AP, 1.3 AP, 1.1 AP and 0.7 AP. We summarize some key characteristics of results: i) improvements are consistent between different backbones, showing the robustness of the proposed UDP method. ii) improvements on small network input method are much more larger than that on

large network input, as larger network input size can effectively suppress the data processing error in state-of-the-art methods.

Results on the test-dev set. Table 2 reports the performance of UDP on COCO test-dev set. The results show better improvement compared with val set, showing superior generalization property of the proposed UDP methods. Specifically, our approach promotes SimpleBaseline by 1.5 AP (70.2 to 71.7) and 1.0 AP (71.9 to 72.9) within ResNet50-256 × 192 and ResNet152-256 × 192 configurations, respectively. For HRNet within W32-256 × 192 and W48-256 × 192 configurations, UDP obtains gains by 1.7 AP (73.5 to 75.2) and 1.4 AP (74.3 to 75.7), respectively. The HRNet-W48-384 × 288 equipped with UDP achieves 76.5 AP and sets a new state-of-the-art for human pose estimation. It is worth noting that our approach only increases negligible calculation burden during training and inference.

Methods	Flip	DT	ED	Shift	DC	AP
A						74.5
B		✓				74.4
C			✓			75.6
D		✓	✓			75.7
E	✓					73.3
F	✓			✓		75.6
G	✓			✓	✓	75.8
H	✓	✓				75.7
I	✓		✓			74.5
J	✓		✓	✓		76.6
K	✓	✓	✓			76.8

Table 3. Ablation study on COCO val set. Flip: Flipping Test. Shift: Shift 1 pixel in post processing [32, 26]. DT: The proposed Data Transformation. ED: The proposed Encoding-Decoding. DC: Direct compensation for ${}^o e(x)'$.

4.3. Ablation Study

In this subsection, we use HRNet-W32 backbone and 256 × 192 input size to perform ablation study. We firstly report the effect of shifting 1 pixel operation used in [26]. As listed in Table 3, method *F* shifts the heatmap of the flipped images by 1 pixel, which brings 2.3 AP improvement from method *E*. And the result obtained by method *E* is even far inferior to the result without using flip strategy (method *A*). This unreported trick (*i.e.* shifting 1 pixel) [32, 6, 26, 16] possibly makes the comparisons unfair with other works.

We then shift \mathbf{k}_a by $-\frac{1}{2s}$ unit length in O_o-X_o direction to directly compensate the existing precision error ${}^o e(x)'$ in [26]. As listed in Table 3, this compensation (method *G*) results in 0.2 AP improvement from method *F*. The improvement is reasonable but small. As analyzed in Section 3.2, the effect of this compensation on error distribution is limited to eliminate the negative effects of data transformation, and the systematic error from encoding-decoding is

dominant.

Next, we only replace the standard data transformation with the proposed data transformation. This modification (method *H* 75.7AP) achieves similar effect as method *G* 75.8 AP which requires shifting 1 pixel and extra compensation for remaining error. The proposed data transformation leads to unbiased results in inference and frees us from complicated post processing (*i.e.* shifting 1 pixel and compensating the remaining error ${}^o e(x)'$). Besides, both method *H* (75.7 AP) and method *G* (75.8 AP) bring a large performance boost when compared with the result from method *E* (73.3 AP), showing the importance of unbiased data transformation for result reproducing.

Meanwhile, we also replace the standard encoding-decoding methods with the proposed ones. This modification achieves 74.5 AP without shifting 1 pixel (method *I*) and 76.6 AP with shifting 1 pixel (method *J*), bringing a performance boost of 1.2 AP improvement and 1.0 AP improvement when compared with the result from standard method [26] *E* and *F*, respectively. The proposed error-free encoding-decoding method shows its superiority for pose estimation problem under fair comparison. The unbiased data transformation provides a precise and fair comparison in method validation which is supposed to be of importance in future research. Method *K* finally brings an impressive performance boost of 1.2 AP by applying the proposed unbiased data processing method on the original HRNet (method *F*).

5. Conclusion and Future Work

In this paper, the common biased data processing for human pose estimation is quantitatively analysed. Interestingly, we find that the systematic errors in standard data transformation and encoding-decoding couple together and significantly degrade the performance of top-down pipelines. Based on analysis results, this paper formulates a principled Unbiased Data Processing (UDP) strategy, which consists of *unit length*-based measurement and a combined classification and regression encoding-decoding approach. The proposed UDP is a model-agnostic strategy and can serve for different top-down pose estimators. In experiments, UDP not only promotes state-of-the-arts on challenging COCO dataset by large margin with variable backbones and input sizes, but also shows its importance in result reproducing and future exploration. Future work will apply the proposed UDP to face landmark, anchor-free object detection and 3D human pose estimation.

References

- [1] Mykhaylo Andriluka, Leonid Pishchulin, Peter Gehler, and Bernt Schiele. 2d human pose estimation: New benchmark and state of the art analysis. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 3686–3693, 2014.

- [2] Mykhaylo Andriluka, Stefan Roth, and Bernt Schiele. Pictorial structures revisited: People detection and articulated pose estimation. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1014–1021, 2009.
- [3] Zhe Cao, Tomas Simon, Shih-En Wei, and Yaser Sheikh. Realtime multi-person 2d pose estimation using part affinity fields. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 7291–7299, 2017.
- [4] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition a new model and the kinetics dataset. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 6299–6308, 2017.
- [5] Kai Chen, Jiangmiao Pang, Jiaqi Wang, Yu Xiong, Xiao Xiao Li, Shuyang Sun, Wansen Feng, Ziwei Liu, Jianping Shi, Wanli Ouyang, et al. Hybrid task cascade for instance segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 4974–4983, 2019.
- [6] Yilun Chen, Zhicheng Wang, Yuxiang Peng, Zhiqiang Zhang, Gang Yu, and Jian Sun. Cascaded pyramid network for multi-person pose estimation. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 7103–7112, 2018.
- [7] Bowen Cheng, Bin Xiao, Jingdong Wang, Honghui Shi, Thomas S Huang, and Lei Zhang. Bottom-up higher-resolution networks for multi-person pose estimation. *arXiv preprint arXiv:1908.10357*, 2019.
- [8] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *IEEE conference on computer vision and pattern recognition*, pages 3213–3223, 2016.
- [9] Hao-Shu Fang, Shuqin Xie, Yu-Wing Tai, and Cewu Lu. Rmpe: Regional multi-person pose estimation. In *IEEE International Conference on Computer Vision*, pages 2334–2343, 2017.
- [10] Pedro F Felzenszwalb, Ross B Girshick, David McAllester, and Deva Ramanan. Object detection with discriminatively trained part-based models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32(9):1627–1645, 2010.
- [11] Kaiming He, Georgia Gkioxari, Piotr Dollr, and Ross Girshick. Mask r-cnn. In *IEEE International Conference on Computer Vision*, 2017.
- [12] Shaoli Huang, Mingming Gong, and Dacheng Tao. A coarse-fine network for keypoint localization. In *IEEE International Conference on Computer Vision*, pages 3028–3037, 2017.
- [13] Eldar Insafutdinov, Leonid Pishchulin, Bjoern Andres, Mykhaylo Andriluka, and Bernt Schiele. Deepcut: A deeper, stronger, and faster multi-person pose estimation model. In *European Conference on Computer Vision*, pages 34–50, 2016.
- [14] Muhammed Kocabas, Salih Karagoz, and Emre Akbas. Multiposenet: Fast multi-person pose estimation using pose residual network. In *European Conference on Computer Vision*, pages 417–433, 2018.
- [15] Peng Li, Jiabin Zhang, Zheng Zhu, Yanwei Li, Lu Jiang, and Guan Huang. State-aware re-identification feature for multi-target multi-camera tracking. In *IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 0–0, 2019.
- [16] Wenbo Li, Zhicheng Wang, Binyi Yin, Qixiang Peng, Yuming Du, Tianzi Xiao, Gang Yu, Hongtao Lu, Yichen Wei, and Jian Sun. Rethinking on multi-stage networks for human pose estimation. *arXiv preprint arXiv:1901.00148*, 2019.
- [17] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European Conference on Computer Vision*, pages 740–755. Springer, 2014.
- [18] Gyeongseok Moon, Ju Yong Chang, and Kyoung Mu Lee. Posefix: Model-agnostic general human pose refinement network. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 7773–7781, 2019.
- [19] Roozbeh Mottaghi, Xianjie Chen, Xiaobai Liu, Nam-Gyu Cho, Seong-Whan Lee, Sanja Fidler, Raquel Urtasun, and Alan Yuille. The role of context for object detection and semantic segmentation in the wild. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 891–898, 2014.
- [20] Alejandro Newell, Zhiao Huang, and Jia Deng. Associative embedding: End-to-end learning for joint detection and grouping. In *Advances in Neural Information Processing Systems*, pages 2277–2287, 2017.
- [21] Alejandro Newell, Kaiyu Yang, and Jia Deng. Stacked hour-glass networks for human pose estimation. In *European Conference on Computer Vision*, pages 483–499. Springer, 2016.
- [22] George Papandreou, Tyler Zhu, Liang-Chieh Chen, Spyros Gidaris, Jonathan Tompson, and Kevin Murphy. Personlab: Person pose estimation and instance segmentation with a bottom-up, part-based, geometric embedding model. In *European Conference on Computer Vision*, pages 269–286, 2018.
- [23] George Papandreou, Tyler Zhu, Nori Kanazawa, Alexander Toshev, Jonathan Tompson, Chris Bregler, and Kevin Murphy. Towards accurate multi-person pose estimation in the wild. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 4903–4911, 2017.
- [24] Leonid Pishchulin, Eldar Insafutdinov, Siyu Tang, Bjoern Andres, Mykhaylo Andriluka, Peter V Gehler, and Bernt Schiele. Deepcut: Joint subset partition and labeling for multi person pose estimation. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 4929–4937, 2016.
- [25] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision*, 115(3):211–252, 2015.
- [26] Ke Sun, Bin Xiao, Dong Liu, and Jingdong Wang. Deep high-resolution representation learning for human pose estimation. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2019.
- [27] Xiao Sun, Bin Xiao, Fangyin Wei, Shuang Liang, and Yichen Wei. Integral human pose regression. In *European Conference on Computer Vision*, pages 529–545, 2018.

- [28] Xiao Sun, Bin Xiao, Fangyin Wei, Shuang Liang, and Yichen Wei. Integral human pose regression. In *European Conference on Computer Vision (ECCV)*, pages 529–545, 2018.
- [29] Jonathan J Tompson, Arjun Jain, Yann LeCun, and Christoph Bregler. Joint training of a convolutional network and a graphical model for human pose estimation. In *Advances in Neural Information Processing Systems*, pages 1799–1807, 2014.
- [30] Alexander Toshev and Christian Szegedy. Deeppose: Human pose estimation via deep neural networks. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1653–1660, 2014.
- [31] Shih-En Wei, Varun Ramakrishna, Takeo Kanade, and Yaser Sheikh. Convolutional pose machines. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 4724–4732, 2016.
- [32] Bin Xiao, Haiping Wu, and Yichen Wei. Simple baselines for human pose estimation and tracking. In *European Conference on Computer Vision*, pages 466–481, 2018.
- [33] Wei Yang, Shuang Li, Wanli Ouyang, Hongsheng Li, and Xiaogang Wang. Learning feature pyramids for human pose estimation. In *IEEE International Conference on Computer Vision*, 2017.
- [34] Wei Yang, Shuang Li, Wanli Ouyang, Hongsheng Li, and Xiaogang Wang. Learning feature pyramids for human pose estimation. In *IEEE International Conference on Computer Vision*, pages 1281–1290, 2017.
- [35] Dongdong Yu, Kai Su, Xin Geng, and Changhu Wang. A context-and-spatial aware network for multi-person pose estimation. *arXiv preprint arXiv:1905.05355*, 2019.
- [36] Jiagang Zhu, Wei Zou, Zheng Zhu, and Yiming Hu. Convolutional relation network for skeleton-based action recognition. *Neurocomputing*, 370:109–117, 2019.
- [37] Jiagang Zhu, Wei Zou, Zheng Zhu, Liang Xu, and Guan Huang. Action machine: Toward person-centric action recognition in videos. *IEEE Signal Processing Letters*, 26(11):1633–1637, 2019.