

DistilPose: Tokenized Pose Regression with Heatmap Distillation

Suhang Ye^{1*} Yingyi Zhang^{2*} Jie Hu^{1*} Liujuan Cao¹

Shengchuan Zhang^{1†} Lei Shen² Jun Wang³ Shouhong Ding² Rongrong Ji¹

¹Key Laboratory of Multimedia Trusted Perception and Efficient Computing, Ministry of Education of China, Xiamen University, ²Tencent Youtu Lab, ³Tencent WeChat Pay Lab³³

Abstract

In the field of human pose estimation, regression-based methods have been dominated in terms of speed, while heatmap-based methods are far ahead in terms of performance. How to take advantage of both schemes remains a challenging problem. In this paper, we propose a novel human pose estimation framework termed DistilPose, which bridges the gaps between heatmap-based and regression-based methods. Specifically, DistilPose maximizes the transfer of knowledge from the teacher model (heatmap-based) to the student model (regression-based) through Token-distilling Encoder (TDE) and Simulated Heatmaps. TDE aligns the feature spaces of heatmap-based and regression-based models by introducing tokenization, while Simulated Heatmaps transfer explicit guidance (distribution and confidence) from teacher heatmaps into student models. Extensive experiments show that the proposed DistilPose can significantly improve the performance of the regression-based models while maintaining efficiency. Specifically, on the MSCOCO validation dataset, DistilPose-S obtains 71.6% mAP with 5.36M parameters, 2.38 GFLOPs, and 40.2 FPS, which saves 12.95 \times , 7.16 \times computational cost and is 4.9 \times faster than its teacher model with only 0.9 points performance drop. Furthermore, DistilPose-L obtains 74.4% mAP on MSCOCO validation dataset, achieving a new state-of-the-art among predominant regression-based models. Code will be available at <https://github.com/yshMars/DistilPose>.

1. Introduction

2D Human Pose Estimation (HPE) aims to detect the anatomical joints of a human in a given image to estimate the poses. HPE is typically used as a preprocessing module that participates in many downstream tasks, such as activity recognition [28, 31], human motion analysis [1], mo-

* Equal contribution. This work was done when Suhang Ye was an intern at Tencent Youtu Lab.

† Corresponding author: zsc_2016@xmu.edu.cn

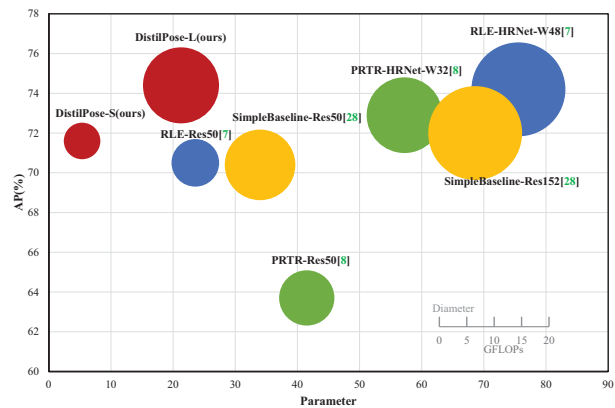


Figure 1. Comparisons between the SOTA methods and the proposed DistilPose on MSCOCO val dataset. Red circles at the upper left corner denote DistilPose. DistilPose outperforms SOTA models in terms of accuracy (AP), Parameter and computational cost (GFLOPs).

tion capture [16], etc. Previous studies on 2D HPE can be mainly divided into two mainstreams: *heatmap-based* and *regression-based* methods. Regression-based methods have significant advantages in speed and are well-suited for mobile devices. However, the insufficient accuracy of regression models will affect the performance of downstream tasks. In contrast, heatmap-based methods can explicitly learn spatial information by estimating likelihood heatmaps, resulting in high accuracy on HPE tasks. But the estimation of likelihood heatmaps requires exceptionally high computational cost, which leads to slow preprocessing operations. Thus, how to take advantages of both heatmap-based and regression-based methods remains a challenging problem.

One possible way to solve the above problem is to transfer the knowledge from heatmap-based to regression-based models [8, 23]. However, due to the different output spaces of regression models and heatmap models (the former is a vector, and the latter is a heatmap), transferring knowledge between heatmaps and vectors faces the following two problems: (1) The regression head usually vectorizes the feature map output by the backbone. And much spatial information will be lost through Global Average Pooling

(GAP) or Flatten operation. Thus, previous work failed to transfer heatmap knowledge to regression models fully. (2) Compared to the coordinate regression, heatmaps naturally contain shape, position, and gradient information [3]. Due to the lack of explicit guidance for such information, regression-based methods are more difficult to learn the implicit relationship between features and keypoints than heatmap-based methods.

In this paper, we propose a novel human pose estimation framework, DistilPose, which learns to transfer the heatmap-based knowledge from a teacher model to a regression-based student model. DistilPose mainly includes the following two components:

(1) A knowledge-transferring module called Token-distilling Encoder (TDE) is designed to align the feature spaces of heatmap-based and regression-based models by introducing tokenization, which consists of a series of transformer encoders. TDE can capture the relationship between keypoints and feature maps/other keypoints [10,32]. (2) We propose to simulate heatmaps to obtain the heatmap information for regression-based students explicitly. The resulting Simulated Heatmaps provide two explicit guidelines, including each keypoint’s 2D distribution and confidence. Note that the proposed Simulated Heatmaps can be inserted between any heatmap-based and regression-based methods to transfer heatmap knowledge to regression models.

DistilPose achieves comparable performance to heatmap-based models with less computational cost and surpasses the state-of-the-art (SOTA) regression-based methods. Specifically, on the MSCOCO validation dataset, DistilPose-S achieves 71.6% mAP with 5.36M parameters, 2.38 GFLOPs and 40.3 FPS. DistilPose-L achieves 74.4% mAP with 21.27M parameters and 10.33 GFLOPs, which outperforms its heatmap-based teacher model in performance, parameters and computational cost. In summary, DistilPose significantly reduces the computation while achieving competitive accuracy, bringing advantages from both heatmap-based and regression-based schemes. As shown in Figure 1, DistilPose outperforms previous SOTA regression-based methods, such as RLE [8] and PRTR [9] with fewer parameters and GFLOPs.

Our contributions are summarized as follows:

- We propose a novel human pose estimation framework, DistilPose, which is the first work to transfer knowledge between heatmap-based and regression-based models losslessly.
- We introduce a novel Token-distilling Encoder (TDE) to take advantage of both heatmap-based and regression-based models. With the proposed TDE, the gap between the output space of heatmaps and coordinate vectors can be facilitated in a tokenized manner.
- We propose Simulated Heatmaps to model explicit

heatmap information, including 2D keypoint distributions and keypoint confidences. With the aid of Simulated Heatmaps, we can transform the regression-based HPE task into a more straightforward learning task that fully exploits local information. Simulated Heatmaps can be applied to any heatmap-based and regression-based models for transferring heatmap knowledge to regression models.

2. Related Work

2.1. Heatmap-based & Regression-based HPE

Heatmap-based pose estimation [10, 17, 21, 22, 25, 29, 30, 32] dominated the area of human pose estimation in terms of performance. Some studies [10, 17, 21, 29, 32] constructed novel networks to extract better features. While others [5, 20, 24, 33] built upon an optimization perspective trying to mitigate quantization errors. In summary, heatmap-based methods made full use of the spatial information of the feature map and obtain a preferable performance. However, efficiency is still a certain drawback of heatmap-based methods.

For regression-based methods, Deeppose [26] is firstly proposed to regress the joint coordinates directly. CenterNet [36] and DirectPose [23] are proposed to accomplish multi-person human pose estimation in a *one-stage* object detection framework, which directly regresses the joint coordinates instead of the bounding box. SPM [18] introduced the root joints to indicate different person instances and hierarchical rooted human body joints representations to better predict long-range displacements for some joints. Recently, RLE [8] introduced a flow model to capture the underlying output distribution and gets a satisfying performance. Although these methods have made great efforts to find the implicit relationship of keypoints, their performance improvement is still insufficient due to the lack of explicit guidance of heatmaps.

2.2. Transformer in HPE

Transformer is proposed in [27] and achieves great success in Natural Language Processing (NLP). Recent studies in vision tasks used Transformer as an alternative backbone to CNN for its ability to capture global dependencies. In the area of 2D human pose estimation, many efforts [9, 10, 14, 15, 19, 30, 32] have been done to incorporate the Transformers. TFpose [14] first introduced Transformer to the pose estimation framework in a regression-based manner. PRTR [9] proposed a two-stage and end-to-end regression-based framework using cascade Transformers and achieves SOTA performance in regression-based methods. TransPose [32] and TokenPose [10] introduced Transformer for heatmap-based human pose estimation achieving comparable performance while being more

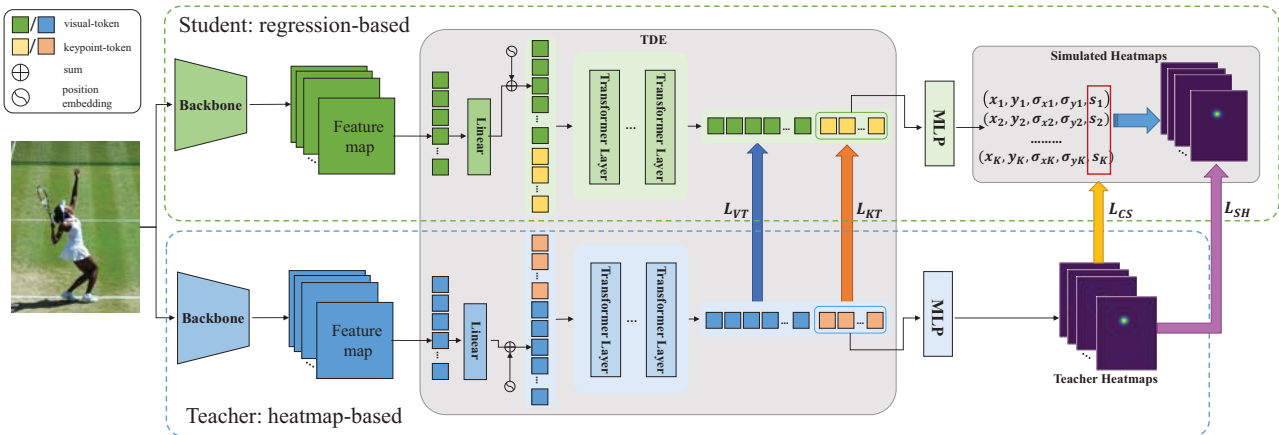


Figure 2. Overall architecture of **DistilPose**. During training, a well-trained and fixed heatmap-based teacher provides its knowledge to help the training of regression-based student with TDE and Simulated Heatmaps.

lightweight. In our work, we introduce the transformer module to assist in finding potential relationships between keypoints.

2.3. Distillation in HPE

Knowledge Distillation (KD) is formally proposed in [4], which aims to transfer the teacher’s learned knowledge to the student model. In 2D human pose estimation, FPD [34] first used knowledge distillation classically based on the Hourglass network. OKDHP [11] introduced an online pose distillation approach that distills the pose structure knowledge in a one-stage manner. ViTPose [30] also implemented a large-to-small model knowledge distillation to prove its knowledge transferability. However, all previous distillation works on human pose estimation ignore the knowledge transferring between heatmap-based and regression-based methods. In this work, for the first time, we propose a heatmap-to-regression distillation framework to take benefits from both schemes.

3. Method

In this section, we propose a distillation-based human pose estimation framework **DistilPose**, the overall framework of which is shown in Fig. 2. In our proposed **DistilPose**, the teacher is a heatmap-based model, and the student is a regression-based model. We transfer the heatmap knowledge of the teacher model to the student model during training and only use the faster student model in the inference stage. **DistilPose** mainly consists of two modules: Token-distilling Encoder and Simulated Heatmaps.

3.1. Token-distilling Encoder

Previous works have tried to introduce the advantages of heatmap models in regression-based methods, such as heatmap pre-training [8], auxiliary heatmap loss [23], etc.

However, the predict-heads cannot be aligned due to the misalignment of the output space. That’s why these works can only perform knowledge transfer on the backbone, which brings models limited performance improvement. According to Fig. 3(a), the heatmap-auxiliary model pays too much attention to regions other than the human body. Hence, we propose a Token-distilling Encoder (TDE) to align the output space of teacher and student by introducing tokenization. By introducing aligned tokenized features, the heatmap knowledge is transferred to the student model losslessly. Thus, the student model learns information that is more focused on the human body itself, as shown in Fig. 3(a).

Specifically, for an input image I , we divide it into several patches according to the size of $pw \times ph$ to form a visual-token. Next, we add K empty nodes as keypoints-token, which are concatenated with visual-token and sent to several transformer encoder layers of TDE. Inspired by LV-ViT [6], we align the visual-tokens and keypoint-tokens between the student and teacher models to obtain the teacher model’s refined attention matrix. As shown in Fig. 3(b), the attention matrix in TDE can learn the relationship between the keypoint-tokens and the visual-tokens of the corresponding position. As for performance improvement, TDE enables our student model to achieve much higher performance than the heatmap-auxiliary training (7.8% \uparrow in Fig. 3(a)).

3.2. Simulated Heatmaps

Basic Distribution Simulation After the head of the student model obtained the aligned knowledge with TDE, we began to think about whether there were other ways to transfer other knowledge of the teacher model further. Existing object detection distillation pipelines give us inspiration [35]: in addition to feature distillation, the output of

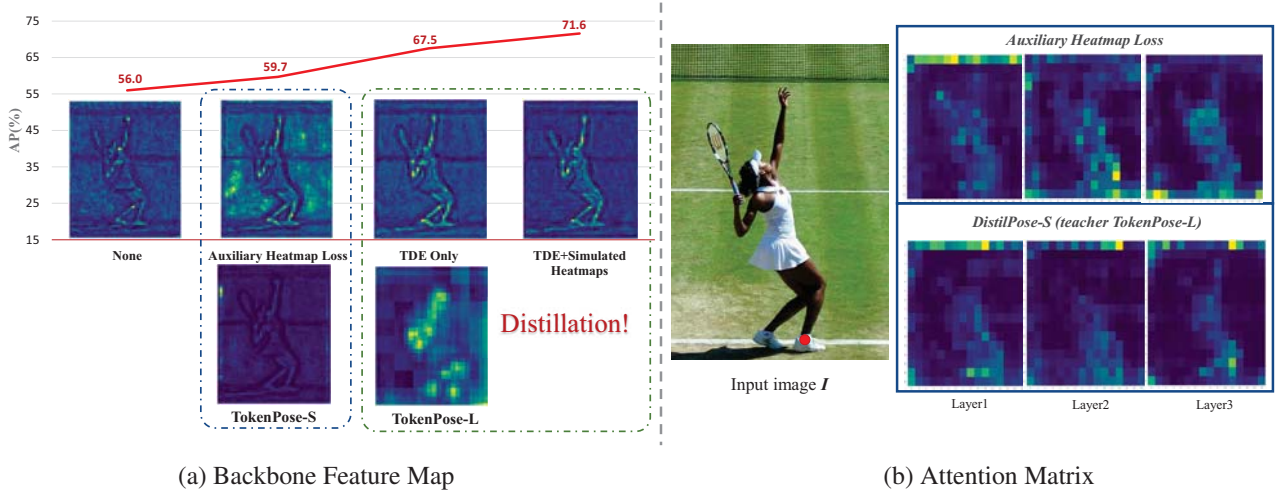


Figure 3. Visualization of backbone feature maps and attention matrix. (a) The first row represent backbone feature maps of direct training, heatmap-auxiliary training, and our proposed distillation training (distill TDE only and distill both TDE and Simulated Heatmaps), respectively. While the second row represent backbone feature maps of the heatmap-based method TokenPose with different backbones. It can be seen here that the full structure of DistilPose is more focused on the human body than the model distill TDE only (the former the feature map background is **darker**). (b) These are attention matrix between the keypoint-token of “right-ankle” (red dot) and visual-tokens.




Figure 4. Visualization of Basis distribution simulation of k^{th} keypoint in Simulated Heatmaps.

the teacher model can also be used as a soft label to further distill knowledge to the student model. Compared with the coordinate vector, the distribution with a well-defined shape and good gradient contained in the heatmap is explicit guide information, which can prompt the model to pay attention to the local information around the keypoints. Not only that, but teacher heatmaps may be closer to reality. It is because there may be erroneous annotations in ground-truth, and the teacher model can filter out these faulty parts by summarizing the knowledge during training. However, heatmaps and coordinate vectors are in different output spaces. How to introduce the distribution information on the heatmap into the space of the coordinate vector has become an urgent problem to be solved.

For this purpose, we do the opposite and propose the concept of Simulated Heatmaps to introduce coordinate vectors into the heatmap space. We let the student model predict both the keypoint coordinates and the correspond-

ing σ values, using a Gaussian distribution to build a virtual heatmap (as shown in Fig. 4). Since the heatmap of the teacher model is not necessarily a regular distribution, we predict σ_x and σ_y from the horizontal/vertical directions, respectively, to construct a heatmap closer to the actual one. The Gaussian distribution is calculated as follows:

$$f(x, y) = \frac{1}{2\pi\sigma_x\sigma_y} e^{-\frac{1}{2}\left(\frac{(x-\mu_x)^2}{\sigma_x^2} + \frac{(y-\mu_y)^2}{\sigma_y^2}\right)}, \quad (1)$$

To align with the ground-truth of the teacher model, we ignore the constant coefficient $\frac{1}{2\pi\sigma_x\sigma_y}$ multiplied by the Gaussian distribution, and the final formula for simulating a heatmap for k_{th} keypoint can be summarized as:

$$H_k(x, y) = e^{-\frac{1}{2}\left(\frac{(x-\mu_{xk})^2}{\sigma_{xk}^2} + \frac{(y-\mu_{yk})^2}{\sigma_{yk}^2}\right)}, \quad (2)$$

where $0 < x \leq hW$, $0 < y \leq hH$ and $0 < k \leq K$. hW and hH are the width and height of the heatmap, K is the total number of keypoints, and H_k refers to the k^{th} simulated heatmap. (μ_{xk}, μ_{yk}) is the predicted keypoints coordinate, and $(\sigma_{xk}, \sigma_{yk})$ is the corresponding deviation pair. In the end, we get K heatmap as shown in Fig. 5.

Confidence Distillation. In addition to distribution information, heatmaps also provide keypoint confidence, which helps filter out incorrect model predictions and is critical in industrial applications. However, our simulated heatmap defaults to 1 at the peak, which cannot be included in the calculation as a valid confidence.

To this end, we prompt the student model to predict the confidence of keypoints s_k directly. Each keypoint confidence s_k corresponds to the keypoint coordinate and is constrained by the corresponding value of the coordinate on the




Figure 5. Visualization of Simulated Heatmaps and teacher heatmaps.

teacher’s heatmap. The keypoint confidence s_k will be accumulated in the human instance score, which is calculated as follows:

$$s^{human} = s^{bbox} * \frac{\sum_{i=1}^K s_k}{K}, \quad (3)$$

where s^{human} is the prediction score of the human instance, and s^{bbox} refers to the confidence score given by the human detector.

3.3. Loss Function

For TDE, we perform distillation for visual-tokens and keypoint-tokens, respectively, and the corresponding loss functions are:

$$\begin{aligned} L_{KT} &= MSE(\mathbf{KT}_h, \mathbf{KT}_r), \\ L_{VT} &= MSE(\mathbf{VT}_h, \mathbf{VT}_r), \end{aligned} \quad (4)$$

where $\mathbf{KT} \in \mathbb{R}^{K \times D}$ and $\mathbf{VT} \in \mathbb{R}^{P \times D}$ refer to keypoint-tokens and visual-tokens. P is the number of patches, and D refers to the dimension of tokens. The subscript h and r refer to the heatmap-based and regression-based model accordingly. MSE refers to the Mean Squared Error Loss used to measure the difference between teacher and student

For the Simulated Heatmaps, we try to align the simulated ones to the teacher heatmaps, and the corresponding loss function is:

$$L_{SH} = \sum_{i=1}^K MSE(H_k, H_{T_k}), \quad (5)$$

where H_{T_k} indicates the k^{th} teacher heatmap.

For the peak constraint of the Simulated Heatmaps, we design the following loss function:

$$L_{CS} = \sum_{i=1}^K |H_{T_k}([x_k], [y_k]) - s_k|, \quad (6)$$

where (x_k, y_k) , $k \in [1, K]$ refers to student coordinate predictions. $[\cdot]$ indicates the round-off operation.

In summary, we train regression-based student models simultaneously under task supervision and knowledge distillation. The overall loss function of our distillation framework is as follows:

$$L = L_{reg} + \alpha_1 L_{KT} + \alpha_2 L_{VT} + \alpha_3 L_{SH} + \alpha_4 L_{CS}. \quad (7)$$

where L_{reg} is the regression task loss for human pose estimation. We use Smooth L1 Loss as L_{reg} to compute the distance between predicted joints and ground-truth coordinates. Invisible keypoints will be filtered, including ground-truth coordinates or predicted coordinates outside the image. In other word, we did not use these invisible keypoints for loss calculation.

4. Experiments

In this section, we evaluate the proposed distillation framework on the MSCOCO dataset. What’s more, we carry out a series of ablation studies to prove the effectiveness and validity of DistilPose. The implementation of our method is based on MMPose [2].

4.1. Implementation Details

4.1.1 Datasets

We mainly conduct our experiments on the MSCOCO dataset [13]. The MSCOCO dataset contains over 200k images and 250k human instances. Each human instance is labeled with $K = 17$ keypoints representing a human pose. Our models are trained on MSCOCO train2017 with 57k images and evaluated on both MSCOCO val2017 and test-dev2017, which contain 5k and 20k images, respectively. Furthermore, the ablation experiments are conducted on the MSCOCO val dataset. We mainly report the commonly used standard evaluation metric Average Precision(AP) as previous works done on the MSCOCO dataset.

4.1.2 Training

We follow the top-down human pose estimation paradigm. All input images are resized into 256×192 resolution. We

Methods	Backbone	Input Size	Param.(M)	GFLOPs	mAP(%)
Heatmap-based Methods					
SimpleBaselines [29]	ResNet-50	256×192	34.0	8.90	70.4
SimpleBaselines [29]	ResNet-101	256×192	53.0	12.40	71.4
SimpleBaselines [29]	ResNet-152	256×192	68.6	15.70	72.0
HRNet [21]	HRNet-W32	256×192	28.5	7.10	74.4
HRNet [21]	HRNet-W48	256×192	63.6	14.60	75.1
TokenPose [10]	stemnet	256×192	6.6	2.40	72.5
TokenPose [10]	HRNet-W48-stage3	256×192	27.5	11.60	75.8
TransPose [32]	ResNet-small	256×192	5.0	5.40	71.5
TransPose [32]	HRNet-Small-W48	256×192	17.5	21.80	75.8
Distillation-based Methods					
OKDHP [12]	2-Stack HG	256×192	13.0	25.50	72.8
OKDHP [12]	4-Stack HG	256×192	24.0	47.00	74.8
Regression-based Methods					
PRTR* [9]	ResNet-50	256×192	41.5	5.45	63.7
PRTR [9]	ResNet-50	384×288	41.5	11.00	68.2
PRTR [9]	ResNet-50	512×384	41.5	18.80	71.0
PRTR* [9]	HRNet-W32	256×192	57.2	10.23	72.9
PRTR [9]	HRNet-W32	384×288	57.2	21.60	73.1
PRTR [9]	HRNet-W32	512×384	57.2	37.80	73.3
RLE [8]	ResNet-50	256×192	23.6	4.04	70.5
RLE* [8]	HRNet-W48	256×192	75.6	15.76	74.2
Poseur [15]	MobileNetV2	256×192	11.36	0.5	71.9
Poseur [15]	ResNet-50	256×192	33.26	4.6	75.4
DistilPose-S (Ours)	stemnet	256×192	5.4	2.38	71.6
DistilPose-L (Ours)	HRNet-W48-stage3	256×192	21.3	10.33	74.4

Table 1. Comparison on MSCOCO *val* dataset. Flip test is used on all methods. * indicates that we re-train and evaluate the models.

Methods	Backbone	Input Size	AP(%)	AP ₅₀ (%)	AP ₇₅ (%)	AP _M (%)	AP _L (%)
PRTR [9]	ResNet-101	384×288	68.8	89.9	76.9	64.7	75.8
PRTR [9]	ResNet-101	512×384	70.6	90.3	78.5	66.2	77.7
RLE* [8]	ResNet-50	256×192	69.8	90.1	77.5	67.2	74.3
DistilPose-S (Ours)	stemnet	256×192	71.0	91.0	78.9	67.5	76.8
PRTR [9]	HRNet-W32	384×288	71.7	90.6	79.6	67.6	78.4
PRTR [9]	HRNet-W32	512×384	72.1	90.4	79.6	68.1	79.0
RLE* [8]	HRNet-W48	256×192	73.7	91.4	81.4	71.1	78.6
DistilPose-L (Ours)	HRNet-W48-stage3	256×192	73.7	91.6	81.1	70.2	79.6

Table 2. Comparison on MSCOCO *test-dev* dataset. * indicates that we re-train and evaluate the models.

adopt a commonly used person detector provided by SimpleBaselines [29] with 56.4% AP for the MSCOCO *val* dataset and 60.9% AP for the MSCOCO *test-dev* dataset. All the models are trained with a batch size of 64 images per GPU, using 8 Tesla V100 GPUs. We adopt Adam as our optimizer and train the models for 300 epochs. The base learning rate is set to $1e-3$, and decays to $1e-4$ and $1e-5$ at the 200^{th} and 260^{th} epoch, respectively. We follow the data augmentation setting in HRNet [21]. We empirically set the hyper-parameters to $\alpha_1 = \alpha_2 = 5e-4$, $\alpha_3 = 1$, $\alpha_4 = 1e-2$.

4.1.3 Model Setting

If not specified, the teacher model we use is a heatmap-based model with a performance of 72.5% (75.2% if extra-post-processing [33] is used during inference) at the cost of 69.41M parameters and 17.03 GFLOPs, which adopts

HRNet-W48 as its backbone.

DistilPose-S adopts a lightweight backbone named *stemnet* from TokenPose [10], which is widely used to down-sample the feature map into $1/4$ input resolution quickly. Stemnet is a very shallow convolutional structure and is trained from scratch. Besides, DistilPose-L uses feature map output by HRNet-W48 [21] at the 3^{rd} stage, following the same setting as TokenPose-L [10].

For the architecture configurations of TDE of DistilPose, the num of transformer layers is 12, embedding dim is 192, the num of heads is 8, and patch size $p_w \times p_h$ is 4×3 .

4.2. Main Results

4.2.1 Comparison with SOTA Regression Methods

We compare the proposed DistilPose with the SOTA regression-based methods on MSCOCO *val* and *test-dev* dataset, and the experimental results are shown in Table 1

Model	Role	Backbone	Methods	Ex-post.	AP(%)	Param(M)	GFLOPs	FPS
TokenPose*	Teacher	HRNet-W48	heatmap	Y	75.2	69.41	17.03	7.8
TokenPose*	Teacher	HRNet-W48	heatmap	N	72.5	69.41	17.03	8.2
DistilPose-S	Student	stemnet	regression	-	71.6 (0.9↓)	5.36 (12.95× ↓)	2.38 (7.16× ↓)	40.2 (4.90× ↑)
DistilPose-L	Student	HRNet-W48-s3	regression	-	74.4 (1.9↑)	21.27 (3.26× ↓)	10.33 (1.65× ↓)	13.7 (1.72× ↑)

Table 3. Comparison with Teacher Model. **Ex-post.** = extra post-processing, which means extra post-processing used to refine the heatmap-to-coordinate transformation during inference. We compute the multiples in comparison with non extra post-processing heatmap-based method. * indicates that we re-train and evaluate the models. And HRNet-W48-s3 is short for HRNet-W48-stage3.

and Table 2, respectively. Also, a line chart is drawn for more visual comparison in Figure 1. We mainly compare our methods with PRTR [9] and RLE [8]. Since PRTR is also Transformer-based methods that achieved SOTA performance among regression-based methods, RLE still dominates regression-based methods. We further compare DistilPose with the latest regression-based SOTA method Poseur [15] in our supplementary materials.

Specifically, DistilPose-S achieves 71.6% with 5.36M and 2.38 GFLOPs, which outperforms PRTR-Res50 by 7.9% at the same input resolution 256×192 while reducing 36.14M (87.1%↓) parameters and 3.07 GFLOPs (56.3%↓). Even if PRTR adopts larger input resolutions (384×288 , 512×384), DistilPose still performs better using 256×192 input resolution.

Compared to non-Transformer SOTA algorithms, DistilPose exceeds RLE in performance, parameter, and computation simultaneously on MSCOCO *val* dataset. As shown in Table 2, we can see that DistilPose performs much better than RLE on large human instances. Since RLE is the first regression-based work with superior performance to contemporary heatmap-based methods and DistilPose outperforms RLE, we also claim that DistilPose achieves comparable performance on par with heatmap-based methods, as shown in Table 1.

4.2.2 Comparison with Teacher Model

We conduct experiments to compare the performance of the student and teacher models in the dimension of AP, Parameter, GFLOPs, and FPS on the MSCOCO *val* dataset. Extra post-processing [5, 33] is always used in the heatmap-based model for the heatmap-to-coordinate transformation during inference to eliminate quantization error, significantly improving performance, but is also followed by extra test-time overhead. In this part, we remove the extra post-processing [33] for more fairly exploring the advantages and disadvantages of heatmap-based and regression-based methods. We report the results between student and teacher without extra post-processing. As shown in Table 3, DistilPose-S sacrifices 0.9% precision but dramatically improves in the reduction of parameter, computation, and test-time overhead. DistilPose-L comprehensively outperformed the heatmap-based teacher. That’s because DistilPose is not only immune to quantization error

Distillation	Simulated Heatmaps		TDE		AP	Improv.
	L_{CS}	L_{SH}	L_{KT}	L_{VT}		
No	-	-	-	-	56.0%	-
Yes	✓				63.2%	+7.2%
		✓			56.4%	+0.4%
	✓	✓			64.1%	+8.1%
			✓		67.1%	+11.1%
				✓	61.7%	+5.7%
	✓	✓	✓	✓	67.5%	+11.5%
	✓	✓	✓	✓	71.6%	+15.6%

Table 4. Ablation studies for different types of knowledge distillation. All types of proposed knowledge transferring benefit the regression-based model, and the combination of all proposed knowledge transferring brings the best performance. **Improv.** = Improvement.

of heatmaps, but also maintains the structural advantages of the regression model [7]. Even if compared to the heatmap-based teacher with extra post-processing, DistilPose-L still achieves comparable performance with much less resource consumption. The experiments above demonstrate that DistilPose can aggregate heatmap and coordinate information and benefit from both schemes.

4.3. Ablation Study

4.3.1 Performance Gain from Different Parts

In this subsection, we conduct several ablation experiments to show how each type of knowledge transfer helps the training of the regression-based student. As shown in Table 4, all types of proposed knowledge transfer benefit the regression-based model. From the perspective of knowledge levels, TDE and Simulated Heatmaps each bring an improvement of 11.5% and 8.1%, respectively. In the simulation heatmap module, the Gaussian simulation heatmaps combined with the confidence prediction can improve the contribution of the two by 0.5%. This proves that confidence predictions and Simulated Heatmaps have mutually reinforcing effects within the model. The combination of all proposed knowledge transfers brings the best performance, which significantly improves the performance by 15.6% comparing to the non-distillation regression-based student model.

Student \ Teacher	Teacher		
	None	stemnet	HRNet-W48
stemnet	56.0%	63.6%	71.6%
HRNet-W48-stage3	63.0%	66.8%	74.4%

Table 5. Ablation studies on different volumes of student and teacher, which demonstrates that regression-based student model performs better with a stronger teacher.

Model	Simulated Heatmaps	Role	mAP	Improv.
SimpleBaseline	-	Teacher	70.4%	-
Deeppose	×	-	52.6%	-
Deeppose	✓	Student	59.7%	+ 6.9%

Table 6. Generalization on CNN-based Model. We only implement Simulated Heatmaps to transfer heatmap-level knowledge on CNN-based models and get a significant performance improvement. **Improv.** = Improvement.

4.3.2 Better Teacher, Better Performance

We conduct experiments to study the performance gain between different students and teachers. As shown in Table 5, we demonstrate that the student models will perform better if they are guided by stronger teacher models. This suggests that the performance of DistilPose we report is not the upper-bound accuracy of our distillation framework but an example to illustrate validity. In practice, we could get help from previous algorithms to train a more potent teacher, which can be further used to enhance the performance of the student model. Or we can simply expand the teacher model’s capacity, finding a trade-off between performance and training memory limitation to maximize the utilization of training resources and not introduce any test-time overhead during inference.

4.3.3 Generalization on CNN-based Model

We also conduct an ablation study on a CNN-based model to show its generalizability. We adopt Deeppose [26] as the student and SimpleBaselines [29] as the teacher. Both are the most concise algorithms in regression-based and heatmap-based models, respectively. We introduce Simulated Heatmaps to transfer heatmap-level knowledge as we do in DistilPose, and ignore TDE for the CNN-based model. Both student and teacher adopt ResNet-50 as the backbone, and the experimental results in Table 6 show that Simulated Heatmaps adapted from DistilPose improve the regression-based student’s performance significantly with no extra test-time overhead.

4.4. Visualization of Confidence Scoring

We show that the confidence score predicted by DistilPose for each coordinate is plausible. As the examples shown in Figure 6, most of the joint predictions are in the right position, and the confidence score predictions of these joints are relatively high. However, there are also some poor

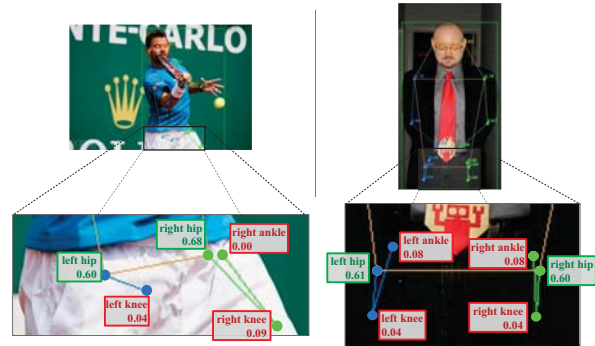


Figure 6. Visualization of confidence scoring. There are 2 rows in each box, representing the type of joint and the confidence score prediction value, respectively. The box border color represents whether the predicted joint is in the correct position (green is correct while red is wrong).

predictions(*e.g.*, the predictions denoted by red box in Figure 6). Fortunately, DistilPose can predict a low confidence score for most poor predictions while keeping high confidence in the correct predictions. Thus we could filter these poor predictions simply by setting a threshold. A plausible confidence score not only enhances the performance of the pose estimator, but is significant in practical applications and downstream tasks for pose estimation.

5. Conclusion

In this work, we proposed a novel human pose estimation framework, termed DistilPose, which includes Token-distilling Module (TDE) and Simulated Heatmaps to perform heatmap-to-regression knowledge distillation. In this way, the student regression model acquired refined heatmap knowledge at both feature and label levels, achieving a big leap in performance while maintaining efficiency. Extensive experiments were conducted on the MSCOCO dataset to demonstrate the effectiveness of our proposed DistilPose. In short, DistilPose achieved state-of-the-art performance among regression-based methods with a much lower computational cost.

Acknowledgement

This work was supported by National Key R&D Program of China (No.2022ZD0118202), the National Science Fund for Distinguished Young Scholars (No.62025603), the National Natural Science Foundation of China (No.U21B2037, No.U22B2051, No.62176222, No.62176223, No.62176226, No.62072386, No.62072387, No.62072389, No.62002305 and No.62272401) and the Natural Science Foundation of Fujian Province of China (No.2021J01002, No.2022J06001). We thank Peng-Tao Jiang from Nankai University for the inspiration during our research.

References

- [1] Lulu Chen, Hong Wei, and James Ferryman. A survey of human motion analysis using depth imagery. *Pattern Recognition Letters*, 34(15):1995–2006, 2013. 1
- [2] MMPose Contributors. Openmmlab pose estimation toolbox and benchmark. <https://github.com/open-mmlab/mmpose>, 2020. 5
- [3] Kerui Gu, Linlin Yang, and Angela Yao. Dive deeper into integral pose regression. In *International Conference on Learning Representations*, 2021. 2
- [4] Geoffrey Hinton, Oriol Vinyals, Jeff Dean, et al. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2(7), 2015. 3
- [5] Junjie Huang, Zheng Zhu, Feng Guo, and Guan Huang. The devil is in the details: Delving into unbiased data processing for human pose estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5700–5709, 2020. 2, 7
- [6] Zi-Hang Jiang, Qibin Hou, Li Yuan, Daquan Zhou, Yujun Shi, Xiaojie Jin, Anran Wang, and Jiashi Feng. All tokens matter: Token labeling for training better vision transformers. *Advances in Neural Information Processing Systems*, 34:18590–18602, 2021. 3
- [7] Haibo Jin, Shengcai Liao, and Ling Shao. Pixel-in-pixel net: Towards efficient facial landmark detection in the wild. *International Journal of Computer Vision*, 129:3174–3194, 2021. 7
- [8] Jiefeng Li, Siyuan Bian, Ailing Zeng, Can Wang, Bo Pang, Wentao Liu, and Cewu Lu. Human pose regression with residual log-likelihood estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11025–11034, 2021. 1, 2, 3, 6, 7
- [9] Ke Li, Shijie Wang, Xiang Zhang, Yifan Xu, Weijian Xu, and Zhuowen Tu. Pose recognition with cascade transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1944–1953, 2021. 2, 6, 7
- [10] Yanjie Li, Shoukui Zhang, Zhicheng Wang, Sen Yang, Wankou Yang, Shu-Tao Xia, and Erjin Zhou. Tokenpose: Learning keypoint tokens for human pose estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11313–11322, 2021. 2, 6
- [11] Zheng Li, Jingwen Ye, Mingli Song, Ying Huang, and Zhigeng Pan. Online knowledge distillation for efficient pose estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11740–11750, 2021. 3
- [12] Zheng Li, Jingwen Ye, Mingli Song, Ying Huang, and Zhigeng Pan. Online knowledge distillation for efficient pose estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 11740–11750, October 2021. 6
- [13] 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. 5
- [14] Weian Mao, Yongtao Ge, Chunhua Shen, Zhi Tian, Xinlong Wang, and Zhibin Wang. Tfpote: Direct human pose estimation with transformers. *arXiv preprint arXiv:2103.15320*, 2021. 2
- [15] Weian Mao, Yongtao Ge, Chunhua Shen, Zhi Tian, Xinlong Wang, Zhibin Wang, and Anton van den Hengel. Poseur: Direct human pose regression with transformers. *arXiv preprint arXiv:2201.07412*, 2022. 2, 6, 7, 11
- [16] Thomas B Moeslund, Adrian Hilton, and Volker Krüger. A survey of advances in vision-based human motion capture and analysis. *Computer vision and image understanding*, 104(2-3):90–126, 2006. 1
- [17] 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. 2
- [18] Xuecheng Nie, Jiashi Feng, Jianfeng Zhang, and Shuicheng Yan. Single-stage multi-person pose machines. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 6951–6960, 2019. 2
- [19] Paschalis Panteleris and Antonis Argyros. Pe-former: Pose estimation transformer. In *International Conference on Pattern Recognition and Artificial Intelligence*, pages 3–14. Springer, 2022. 2
- [20] 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 *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4903–4911, 2017. 2
- [21] Ke Sun, Bin Xiao, Dong Liu, and Jingdong Wang. Deep high-resolution representation learning for human pose estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5693–5703, 2019. 2, 6
- [22] Xiao Sun, Bin Xiao, Fangyin Wei, Shuang Liang, and Yichen Wei. Integral human pose regression. In *Proceedings of the European Conference on Computer Vision (ECCV)*, September 2018. 2
- [23] Zhi Tian, Hao Chen, and Chunhua Shen. Directpose: Direct end-to-end multi-person pose estimation. *arXiv preprint arXiv:1911.07451*, 2019. 1, 2, 3
- [24] Jonathan Tompson, Ross Goroshin, Arjun Jain, Yann LeCun, and Christoph Bregler. Efficient object localization using convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 648–656, 2015. 2
- [25] Jonathan J Tompson, Arjun Jain, Yann LeCun, and Christoph Bregler. Joint training of a convolutional network and a graphical model for human pose estimation. *Advances in neural information processing systems*, 27, 2014. 2
- [26] Alexander Toshev and Christian Szegedy. Deeppose: Human pose estimation via deep neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1653–1660, 2014. 2, 8
- [27] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. 2

- [28] Arpita Vats and David C. Anastasiu. Key point-based driver activity recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pages 3274–3281, June 2022. 1
- [29] Bin Xiao, Haiping Wu, and Yichen Wei. Simple baselines for human pose estimation and tracking. In *Proceedings of the European Conference on Computer Vision (ECCV)*, September 2018. 2, 6, 8
- [30] Yufei Xu, Jing Zhang, Qiming Zhang, and Dacheng Tao. Vit-pose: Simple vision transformer baselines for human pose estimation. *arXiv preprint arXiv:2204.12484*, 2022. 2, 3
- [31] Santosh Kumar Yadav, Achleshwar Luthra, Kamlesh Tiwari, Hari Mohan Pandey, and Shaik Ali Akbar. Arfdnet: An efficient activity recognition & fall detection system using latent feature pooling. *Knowledge-Based Systems*, 239:107948, 2022. 1
- [32] Sen Yang, Zhibin Quan, Mu Nie, and Wankou Yang. Transpose: Keypoint localization via transformer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11802–11812, 2021. 2, 6
- [33] Feng Zhang, Xiatian Zhu, Hanbin Dai, Mao Ye, and Ce Zhu. Distribution-aware coordinate representation for human pose estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 7093–7102, 2020. 2, 6, 7
- [34] Feng Zhang, Xiatian Zhu, and Mao Ye. Fast human pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3517–3526, 2019. 3
- [35] Zhaohui Zheng, Rongguang Ye, Ping Wang, Dongwei Ren, Wangmeng Zuo, Qibin Hou, and Ming-Ming Cheng. Localization distillation for dense object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9407–9416, 2022. 3
- [36] Xingyi Zhou, Dequan Wang, and Philipp Krähenbühl. Objects as points. *arXiv preprint arXiv:1904.07850*, 2019. 2