

Self-Correctable and Adaptable Inference for Generalizable Human Pose Estimation

Zhehan Kan¹, Shuoshuo Chen¹, Ce Zhang¹, Yushun Tang¹, Zhihai He^{1,2*}

¹Department of Electronic and Electrical Engineering, Southern University of Science and Technology, Shenzhen, China

²Pengcheng Laboratory, Shenzhen, China

{kanzh2021, chenss2021, zhangc2019, tangys2022}@mail.sustech.edu.cn, hezh@sustech.edu.cn

Abstract

A central challenge in human pose estimation, as well as in many other machine learning and prediction tasks, is the generalization problem. The learned network does not have the capability to characterize the prediction error, generate feedback information from the test sample, and correct the prediction error on the fly for each individual test sample, which results in degraded performance in generalization. In this work, we introduce a self-correctable and adaptable inference (SCAI) method to address the generalization challenge of network prediction and use human pose estimation as an example to demonstrate its effectiveness and performance. We learn a correction network to correct the prediction result conditioned by a fitness feedback error. This feedback error is generated by a learned fitness feedback network which maps the prediction result to the original input domain and compares it against the original input. Interestingly, we find that this self-referential feedback error is highly correlated with the actual prediction error. This strong correlation suggests that we can use this error as feedback to guide the correction process. It can be also used as a loss function to quickly adapt and optimize the correction network during the inference process. Our extensive experimental results on human pose estimation demonstrate that the proposed SCAI method is able to significantly improve the generalization capability and performance of human pose estimation.

1. Introduction

Human pose estimation (HPE) aims to correctly predict and localize human body joints. A variety of downstream applications are based on human pose estimation, such as motion capture [7, 27], activity recognition [1, 6, 37], person tracking [36, 41] and video surveillance [18]. Recently, deep learning-based methods for human pose estimation

have achieved remarkable success [2, 4, 12, 26, 28, 30]. However, in complex or unseen scenarios, pose estimation remains very challenging due to occlusions, cluttered background, and large variations of appearance and scenes, especially for those distal keypoints at the end locations of body parts, such as wrists and ankles, which have large degrees of motion freedom and often suffer from severe occlusions [13, 39].

We recognize that one major challenge in current human pose estimation, as well as in many other prediction tasks, is generalization. Network models, which have been well learned on the training set, often experience significant performance degradation on the test samples which are collected from different environments or scenarios. For example, in human pose estimation, there are different types of occlusions of body parts due to complex scene structures and free-style motions of human bodies. More importantly, the occlusion scenarios of the test samples could be much different from those in the training samples. This often leads to the significant performance degradation of human pose estimation from the training data to the test data. For example, in our experiment, the average prediction accuracy on the training samples is 95.5%. However, on the test set, this accuracy drops to 67%. For those distal keypoints at tip locations of body parts which often experience more significant occlusions, their average performance drop is much more significant, from 95.3% to 57%.

To address this performance degradation or generalization problem, there are two major questions that need to be carefully answered: (1) how can we tell if the prediction is accurate or not during testing and how to characterize the prediction error? This is difficult because the ground truth values of the test samples are not available during testing. Specifically, in pose estimation, we do not have the labeled ground truth locations of the body keypoints. (2) How to correct the prediction error based on the specific characteristics of the test sample? Current network models, once successfully trained with labeled samples at the training side, remain fixed during testing, performing the feed-

*Corresponding author.

forward-only inference process to generate the prediction result. There is no mechanism for us to examine the specific characteristics of the test sample and use them as feedback to correct the prediction error or adjust the network model. We believe that this unique capability of sample-specific prediction error characterization, error correction, and model optimization is very important for the generalization performance of learned network models. It also has the potential to significantly improve the prediction accuracy of test samples.

To address these two challenging issues, in this work, we propose to explore a learning-based feedback-control or correction method for prediction, with applications to human pose prediction. Specifically, let $\hat{\mathbf{v}} = \Phi(\mathbf{u})$ be the prediction network which is tasked to predict the true value of \mathbf{v} from input \mathbf{u} . To answer the first question, we design and learn a fitness feedback network Γ which compares the prediction result $\hat{\mathbf{v}} = \Phi(\mathbf{u})$ of the prediction network Φ against the original input \mathbf{u} and generate a self-referential feedback error. Very interestingly, in this work, we find that this self-referential feedback error is highly correlated with the prediction error of the network Φ . Note that, when computing the self-referential error, we do not need the ground truth data. It can be directly computed on the input sample using the prediction-feedback networks. This allows us to characterize the prediction error of test samples.

Under the guidance of self-referential error feedback, we train a prediction error correction network \mathbf{C} to adjust the inference results during the prediction process to improve the prediction accuracy for the test samples. Besides, we find that the self-referential error and the fitness feedback network (FFN) can be used to construct a self-referential loss function on the test samples to quickly adapt and optimize the network model during the inference stage, making the model learnable on the test side. We apply the above self-correctable and adaptable inference (SCAI) method to human pose estimation. Our extensive experimental results on benchmark datasets demonstrate that the proposed SCAI method is able to significantly improve the generalization capability of the underlying prediction algorithm. It outperforms the existing state-of-the-art methods on human pose estimation by large margins. For example, on the MS COCO-testdev dataset, our method improves upon the current best method by up to 1.4%, which is quite significant.

2. Related Work

In this section, we review related works on human pose estimation, pose refinement, network generalization, and model adaptation.

(1) Human pose estimation. There are two major categories of methods developed for human pose estimation: *top-down* and *bottom-up* approaches. *Top-down* methods [4, 13, 20, 28, 31] first detect all persons in an image and

then predict keypoint locations for each person. For example, Chen *et al.* [4] utilized a human detector to generate bounding boxes as the inputs into a feature pyramid network for keypoint estimation. *Bottom-up* methods [19, 24, 40] first detect the joints for all persons in the input image and then group them into individual persons. For example, Li *et al.* [19] grouped detected keypoints into instances based on a greedy assignment algorithm. Luo *et al.* [24] introduced scale-adaptive and weight-adaptive heatmap regression to alleviate large differences in human scales.

(2) Pose refinement. A number of methods have been developed for pose refinement [3, 9, 25]. Fieraru *et al.* [9] refined the estimation by exploiting dependencies between the input image and the body structure of the human pose through a network. Moon *et al.* [25] trained a refinement network with synthetic poses generated from error distribution as input. Wang *et al.* [35] extracted guidance keypoints from coarse pose estimation and applied a GCN as a refinement module to seek mutual information among keypoints. It should be noted that these refinement methods did not consider the sample distribution shift between the train set and test set and have not addressed the generalization problem effectively. Kan *et al.* [15] developed a prediction-verification network and perform a local search within the neighborhood of the prediction result. This multi-iteration search method is very time-consuming.

(3) Generalization and adaptation. Generalization is a significant challenge for existing network learning methods. Test-time adaptation [32, 34] aims to improve generalization ability of the deep learning models when the test set exhibits a different data distribution from that of the train set. Sun *et al.* [32] firstly trained self-supervised tasks on the test samples with model parameters optimized and then performed inference. Tung *et al.* [33] performed test-time optimization with self-supervised losses driven by re-projection errors of keypoint, segmentation, and motion against detected 2D versions respectively. Li *et al.* [22] constructed a transformation between the self-supervised and the supervised keypoints, so the model can be fine-tuned toward the self-supervised objective of image reconstruction during inference. However, the self-supervised objectives of test-time adaptation utilized by these methods require specific conditions, such as projection from 3D to 2D or a set of images of the same person. Our method can predict human pose given general 2D images. Besides, our proposed self-referential error is highly correlated with the pose prediction error, which can be used to guide the prediction correction process and to construct self-referential loss for network model adaptation during test time.

3. Method

We first present our SCAI method to address the generalization challenge in generic network learning and predic-

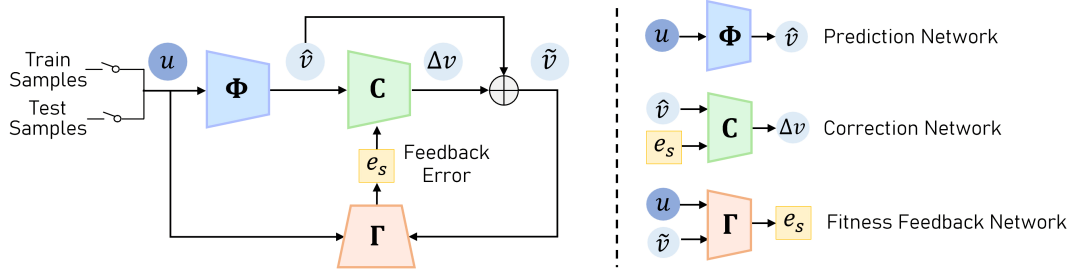


Figure 1. An overview of our self-correctable and adaptable inference (SCAI) method.

tion. We then explain how this method can be applied to the specific problem of human pose estimation.

3.1. Self-Correctable and Adaptable Inference

In this section, we present the proposed self-correctable inference and adaptable inference method.

(1) Overview. As illustrated in Figure 1, let $\hat{v} = \Phi(\mathbf{u})$ be the prediction network which is tasked to predict the true value of \mathbf{v} from input \mathbf{u} . We learn a prediction error correction network \mathbf{C} to produce a correction $\Delta\mathbf{v}$ which adjusts the prediction results \hat{v} to $\tilde{v} = \hat{v} + \Delta\mathbf{v}$, aiming to improve the prediction accuracy on the test samples. To guide such correction, we design an FFN Γ with the original input \mathbf{u} and the corrected prediction \tilde{v} as inputs to generate a self-referential error. This error is used as important feedback to guide the correction network \mathbf{C} to perform adaptive correction of the predicted result. Besides, we find that the self-referential feedback error can be also used to define a loss function to quickly adapt and optimize the network model during the inference stage, making the model learnable on the test side.

(2) Self-Correctable Inference. A deep neural network Φ , already being fully optimized on the training set, often suffers from significant performance degradation when applied to new or unseen test samples. The prediction model remains fixed during testing, performing the feed-forward-only inference process to generate the vanilla prediction results $\hat{v} = \Phi(\mathbf{u})$. As discussed above, we propose to learn a correction network \mathbf{C} to correct the prediction error and improve the prediction results from \hat{v} to \tilde{v} conditioned by an error feedback e_s

$$\Delta\mathbf{v} = \mathbf{C}(\hat{v}|e_s), \quad \tilde{v} = \hat{v} + \Delta\mathbf{v}. \quad (1)$$

Note that the conditional error feedback e_s is very important for the correction network. Without this feedback, the correction network cannot achieve any performance improvement since the original prediction network Φ is already well-trained and fully optimized. Adding another feed-forward correction network simply does not help.

Now, the question is how can we generate a useful feedback signal e_s for the correction network \mathbf{C} ? To address

this issue, we propose to introduce a fitness feedback network Γ . It aims to evaluate how good the corrected prediction result is. Note that we do not know the ground truth of the prediction result. Our idea is to map the prediction result to the original input domain and compare it against the original input which has its ground truth value. Specifically, the fitness feedback network Γ takes two inputs, the corrected prediction \tilde{v} and the original input \mathbf{u} . The output of Γ is the so-called *self-referential feedback error* e_s :

$$e_s = \Gamma(\tilde{v}, \mathbf{u}). \quad (2)$$

which is used to guide the correction network \mathbf{C} . From Section 3.2, specifically, Figure 3, we will see that, with successful training, e_s is highly correlated with the prediction error. This strong correlation allows us to use e_s as feedback to guide the correction of the prediction result. Otherwise, if the correction is weak, the correction process becomes unreliable and cannot achieve improved prediction performance.

(3) Self-Adaptable Inference. As its unique feature, the feedback error is self-referential. In other words, when computing this error, we do not need the ground truth value of the prediction. We only need the prediction network Φ , the fitness feedback network Γ , the correction network \mathbf{C} , and the input. This implies that we can also compute this feedback error e_s during the network inference process. Once computed, its norm can be used as a loss function to quickly adapt and optimize the correction network \mathbf{C} model during the inference stage using gradient back-propagation. During the update, the prediction network Φ and the FFN Γ remain fixed. It should be noted that this adaptation and optimization is only for the current test sample. Certainly, it can be extended to a batch or a cluster of test samples to reduce the complexity. In the following section, we will use human pose estimation as an example to explain the specific training procedures and loss function design for the proposed SCAI method.

3.2. SCAI for Human Pose Estimation

In this work, we use human pose estimation as an example to implement the proposed SCAI method and demon-

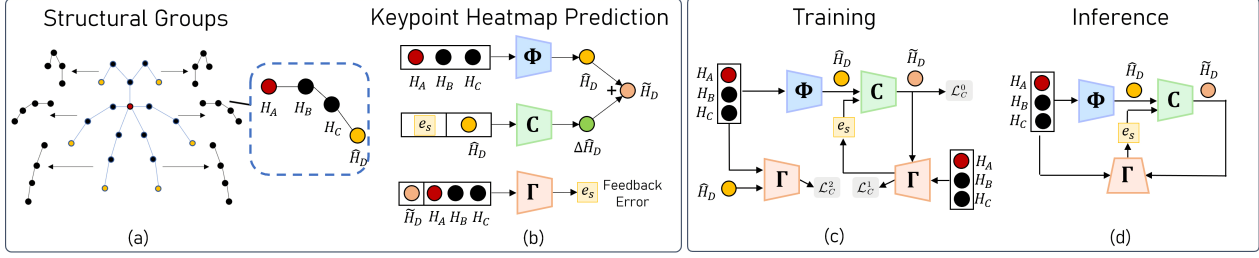


Figure 2. (a) Structural grouping of body keypoints. (b) Main network models. (c) Training process and (d) inference process of our SCAI method for human pose estimation. Only the refinement of the distal keypoints heatmap \hat{H}_D is depicted as an example.

strate its performance. The major question is: how should we re-structure the pose estimation problem so that we can apply the proposed prediction-feedback-correction scheme?

(1) Structural groups of body keypoints. Human pose estimation refers to the detection of body keypoint locations based on an input RGB image I with dimensions $W \times H \times 3$. The task involves identifying K keypoints $X = \{X_1, X_2, \dots, X_K\}$ with high precision. Heatmap-based methods transform this problem into estimating K heatmaps $\{H_1, H_2, \dots, H_K\}$ of size $W' \times H'$. The location of a keypoint can be determined using different methods, such as grouping or peak finding, based on the corresponding heatmap [39]. For instance, the pixel with the highest heatmap value can be selected as the keypoint location.

To enhance the generalization and prediction performance of human pose estimation using the proposed SCAI method, we adopt the approach introduced by Kan *et al.* [15] to partition the body keypoints into six structural groups, as depicted in Figure 2(a). Each structural group corresponds to a body part, with keypoints that are connected during motion. The group is further divided into a distal keypoint X_D at the tip location of the body part, such as the wrists and ankles, and proximal keypoints X_A, X_B, X_C . We observe that distal keypoints often have larger errors in prediction due to more significant freedom of motion and possible occlusion by other objects. Our main objective is to leverage the proposed SCAI approach to improve the prediction accuracy and generalization capability of these distal keypoints.

(2) SCAI Network Design. Given a set of keypoint heatmaps $\{H_A, H_B, H_C, H_D\}$ estimated by a baseline model, *e.g.* HRNet [30], our task is to refine heatmaps $\{H_B, H_C, H_D\}$. Let us consider the refinement of heatmap H_D of the distal keypoint X_D as an example. The input to the prediction network Φ is $\mathbf{u} = \{H_A, H_B, H_C\}$, the heatmaps for the proximal keypoints. The prediction output is $\hat{\mathbf{v}} = \hat{H}_D$. The correction network \mathbf{C} is used to refine the prediction result into $\tilde{\mathbf{v}} = \tilde{H}_D$. The output of FFN Γ is the feedback error \mathbf{e}_s to guide the correction.

As illustrated in Figure 2(b), the prediction network Φ predicts the distal keypoint heatmap $\hat{H}_D = \Phi(H_A, H_B, H_C)$ using heatmaps of proximal keypoints

$\{H_A, H_B, H_C\}$. The correction network \mathbf{C} aims to generate a corrected and improved estimation of H_D conditioned by the feedback error \mathbf{e}_s

$$\tilde{H}_D = \hat{H}_D + \mathbf{C}(\hat{H}_D | \mathbf{e}_s). \quad (3)$$

Here, \mathbf{e}_s is generated by the fitness feedback network Γ , which has two inputs, the corrected prediction result \tilde{H}_D and the original inputs $[H_A, H_B, H_C]$:

$$\mathbf{e}_s = \Gamma([H_A, H_B, H_C], \tilde{H}_D). \quad (4)$$

In the following experiment, we demonstrate that the L_2 -norm of self-referential feedback error \mathbf{e}_s is highly correlated with the actual network prediction error. In this human pose estimation experiment, we choose 400 batches of 25600 test samples to show this correlation. In Figure 3, the vertical axis shows the average prediction accuracy of the body keypoints in each batch of test samples. The horizontal axis shows their average self-referential error. Each dot represents one test batch. We can see that there is a very strong correlation between them, the corresponding correlation coefficient is -0.84 .

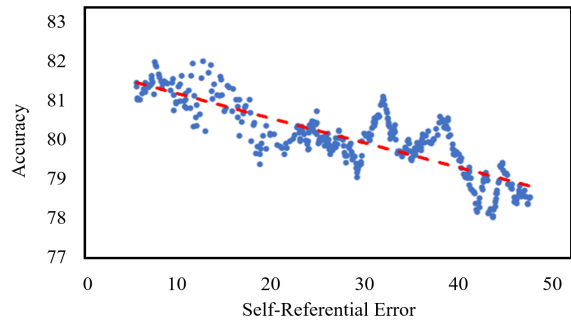


Figure 3. Strong correlation between self-referential error and network prediction accuracy.

(3) SCAI Network Training. In the following, we explain how these three networks are trained with labeled samples. Note that, at the training side, all the predicted keypoints have their ground truth values $\{H_A^*, H_B^*, H_C^*, H_D^*\}$. Therefore, for the prediction network Φ , its loss function is given by $\mathcal{L}_\Phi = \|\hat{H}_D - H_D^*\|_2$, which

is the L_2 distance between the predicted heatmap and the ground truth value [30]. For the fitness feedback network Γ , it maps the prediction result \tilde{H}_D to the original input domain $\{H_A, H_B, H_C\}$ and compare it against the original input. So, during the training stage, its loss function can be expressed in the form of self-referential error, given by $\mathcal{L}_\Gamma = \|\hat{H}_A - H_A^*\|_2$. For the correction network C , its loss function is given by

$$\mathcal{L}_C = a \cdot \mathcal{L}_C^0 + b \cdot \mathcal{L}_C^1 + \lambda \cdot (\mathcal{L}_C^1 - \mathcal{L}_C^2), \quad (5)$$

where $\mathcal{L}_C^0 = \|\tilde{H}_D - H_D^*\|_2$, $\mathcal{L}_C^1 = \|\hat{H}_A - H_A^*\|_2$ and $\mathcal{L}_C^2 = \|\bar{H}_A - H_A^*\|_2$, with a , b and λ representing the weights for the three losses, respectively. \mathcal{L}_C^1 is defined by the distance between the mapping from the corrected prediction \tilde{H}_D and the original input. \mathcal{L}_C^2 is defined by the distance between the mapping from the uncorrected prediction \bar{H}_D and the original input. Here, the last term $\mathcal{L}_C^1 - \mathcal{L}_C^2$ is used to ensure the effectiveness of correction: the corrected result has a smaller self-referential error than the original prediction. Specifically, if the correction network improves the accuracy of keypoint prediction, the self-referential error evaluated by FFN will be smaller than the one without correction, that is, \mathcal{L}_C^1 value is smaller than \mathcal{L}_C^2 . Thus, the loss term $\mathcal{L}_C^1 - \mathcal{L}_C^2$ is able to guide the training of the correction network to effectively refine the keypoint prediction.

During the training process, the prediction network Φ is pre-trained using the training samples $\{(H_A, H_B, H_C) \rightarrow H_D\}$. The FFN is also pre-trained with training samples $\{(H_B, H_C, H_D) \rightarrow H_A\}$. Here, \rightarrow represents the network prediction. During the training stage, the prediction network Φ is fixed. The pre-trained model of the FFN Γ is used as its initial model. The FFN and the correction networks are then jointly trained using their loss functions. Similar to the refinement of proximal keypoint X_D , we can also develop correction networks for keypoints X_B and X_C to optimize their heatmaps H_B and H_C , respectively. More training details are shown in Supplemental Materials.

(4) Self-referential Adaptable Inference for Human Pose Estimation. Note that the self-referential error is computed based on the input samples and does not require the ground truth data of any body keypoints. Therefore, on the test side, we use this self-referential error as a loss function to update the network model. Specifically, in this work, we choose to update the correction network C while other networks, including the prediction network Φ and the FFN Γ remain fixed. It should be noted that this model refinement is performed separately for each batch of test images. When moving to a new batch, the initial models obtained from the training set are restored and then refined. In other words, the models learned from one batch are not used for the next test batch to ensure flexible model adaptation. Figure 4(a) shows the decreasing self-referential error and the convergence behavior of this model learning and adapta-

tion process. Figure 4(b) shows that the accuracy of the test batch is consistently improved with the training epochs. This demonstrates that our self-referential adaptable inference method is able to use the test samples as feedback to update the network models and improve the prediction performance, providing natural and enhanced generalization capability for the learning and prediction network.




Figure 4. The figure on the left is the loss decay curve of SCAI during training in the inference process, and the figure on the right is the test result corresponding to each epoch.

4. Experiments

This section presents a thorough evaluation of the proposed SCAI method, which includes experimental results, performance comparisons with state-of-the-art methods, and ablation studies.

4.1. Datasets

In this section, we present the performance comparisons and ablation studies of our proposed SCAI method on two challenging datasets, namely the MS COCO [23] and CrowdPose [20]. The MS COCO dataset is a widely-used benchmark for human pose estimation, comprising 64K images with 270K annotated persons and 17 keypoints. The dataset includes diverse poses of multiple persons with varying body scales and occlusion patterns, making it challenging for human pose estimation models. Our models are trained on the train2017 split, which consists of 57K images with 150K persons, and ablation studies are conducted on the val2017 split. The CrowdPose dataset contains 20K images with 80K annotated persons and 14 keypoints. It includes crowded scenes and poses with various challenges. Unlike MS COCO, we partition the keypoints into 4 groups instead of 6 groups. We train our models on the train set, which has 10K images and 35.4K persons and evaluate them on the validation set (2K images with 8K persons) and the test set (8K images with 29K persons).

4.2. Experimental Settings

For fair comparisons, we use the HRNet and ResNet as our baseline, from which pre-predicted results are produced to be further refined by our method. We follow the same training configuration in existing works [30, 39]. The correction network is trained with a full convolution network.

Table 1. Comparison with the state-of-the-art methods on COCO test-dev. The best results are in **bold** and the second best results are underlined.

Method	Backbone	Size	AP	AP^{50}	AP^{75}	AP^M	AP^L	AR
CFN [14]	-	-	72.6	86.1	69.7	78.3	64.1	-
CPN(ensemble) [4]	ResNet-Incep.	384×288	73.0	91.7	80.9	69.5	78.1	79.0
CSM+SCARB [28]	R152	384×288	74.3	91.8	81.9	70.7	80.2	80.5
CSANet [42]	R152	384×288	74.5	91.7	82.1	71.2	80.2	80.7
HRNet [30]	HR48	384×288	75.5	92.5	83.3	71.9	81.5	80.5
MSPN [21]	MSPN	384×288	76.1	93.4	83.8	72.3	81.5	81.6
PoseFix [25]	HR48+R152	384×288	76.7	92.6	84.1	73.1	82.6	81.5
DARK [43]	HR48	384×288	76.2	92.5	83.6	72.5	82.4	81.1
UDP [13]	HR48	384×288	76.5	92.7	84.0	73.0	82.4	<u>81.6</u>
Graph-PCNN [35]	HR48	384×288	76.8	92.6	84.3	73.3	82.7	<u>81.6</u>
SCIO [15]	HR48	384×288	<u>79.2</u>	<u>93.5</u>	<u>85.8</u>	74.1	<u>84.2</u>	<u>81.6</u>
SCAI (Ours)	HR48	384×288	80.6	94.8	87.0	<u>78.1</u>	84.8	83.1
Performance Gain			+1.4	+1.3	+1.2	-0.2	+0.6	+1.5

We use the Adam [17] optimizer for training. More experimental details are provided in the Supplemental Materials.

Table 2. Comparison with the state-of-the-art methods on CrowdPose test-dev. The best results are in **bold** and the second best results are underlined.

Method	Backbone	AP	AP^{med}
Mask-RCNN [12]	ResNet101	60.3	-
AlphaPose [8]	-	61.0	61.4
OccNet [10]	ResNet50	65.5	66.6
JC-SPPE [20]	ResNet101	66.0	66.3
HigherHRNet [5]	HR48	67.6	-
MIPNet [16]	HR48	70.0	71.1
SCIO [15]	HR48	<u>71.5</u>	<u>72.2</u>
SCAI (Ours)	HR48	72.4	73.2
Performance Gain		+0.9	+1.0

Table 3. Comparison with state-of-the-art of three baselines on COCO test-dev.

Method	Backbone	AP	AP^M	AR
SimpleBaseline	R152	73.7	70.3	79.0
+ SCAI	R152	78.8	75.5	82.5
Performance Gain		+5.1	+5.2	+3.5
HRNet	HR32	74.9	71.3	80.1
+ SCAI	HR32	79.9	76.8	82.5
Performance Gain		+5.0	+5.5	+2.4
HRNet	HR48	75.5	71.9	80.5
+ SCAI	HR48	80.6	78.1	83.1
Performance Gain		+5.1	+6.2	+2.6

4.3. Performance Comparisons

In Table 1, we compare the performance of our SCAI method with the following state-of-the-art methods on the COCO test set: the CFN [14], CPN (ensemble) [4], CSM+SCARB [28], CSANet [42], HRNet [30], MSPN [21], PoseFix [25], DARK [43], UDP [13], Graph-PCNN [35] and SCIO [15] methods. We can see that our SCAI method outperforms the current state-of-the-art methods by large margins, up to 1.4%.

Table 2 presents the results of our SCAI method compared to other state-of-the-art methods on the challenging CrowdPose dataset. The methods compared include Mask-RCNN [12], SimpleBaseline [39], AlphaPose [8], OccNet [10], JC-SPPE [20], HigherHRNet [5], MIPNet [16], and SCIO [15]. Our SCAI method achieves an improvement in average precision by up to 0.9% compared to the current best method SCIO [15]. This indicates that our method can provide more accurate pose estimation for multi-person scenes with challenging occlusion scenarios.

Table 3 provides a comparison of our SCAI method with other state-of-the-art methods that use different baseline networks such as R152, HR32, and HR48 networks. Our SCAI method consistently outperforms other methods, even when different baseline networks are employed.

As Figure 5 shows, compared to the local search proposed in SCIO [15], our method converges much faster and is able to stably correct the heatmap to be more accurate during the optimization process.

4.4. Ablation Studies

To assess the efficacy of our proposed SCAI method and analyze the influence of individual algorithmic components, we performed various ablation experiments on the COCO test-dev dataset. Our method incorporates two key inno-




Figure 5. Comparison between the local search and the correction from SCIO and SCAI, where t represents the number of iterations.

variations, namely self-correctable inference (SCI) and self-adaptable inference (SAI). Self-correctable inference consists of three essential parts: the correction network, self-referential error, and the joint training of the correction network and FFN. In Table 4, we present the results of our ablation studies. The first row shows the performance of the baseline method. The second row indicates the performance when incorporating the correction network, while the third and fourth rows correspond to the addition of the self-referential error and joint training components, respectively. The last row reports the results when including the self-adaptable inference component. Our findings demonstrate that each algorithmic component makes a substantial contribution to the overall performance of our method.

Table 4. Ablation study on different algorithmic components on COCO test-dev dataset.

Method	AP	AP^M	AR
Baseline	75.5	71.9	80.5
Baseline + SCI			
+ Correction Network	78.5	73.9	81.4
+ Self-Referential Error	78.9	75.8	81.6
+ Joint Training	79.8	77.8	82.6
Baseline + SCI + SAI			
+ Self-Adaptable Inference	80.6	78.1	83.1

Table 5 presents an analysis of the impact of varying loss weights in Equation (5) on our experimental results. The results suggest that the optimal weights for the parameters a , b , and λ are 0.85, 0.65, and 0.45, respectively.

Table 5. Ablation study on different loss parts and their weights in Equation (5).

a	b	λ	AP	AP^M	AR
0.85	0.65	0.30	80.2	77.6	82.7
0.85	0.65	0.60	80.4	77.9	82.9
0.85	0.65	0.45	80.6	78.1	83.1
1.00	0.65	0.45	80.5	78.1	83.1
0.85	0.40	0.45	80.4	77.9	83.1

Figure 6 presents a comparison of the pose estimation results refined by our proposed SCAI method against those produced by the baseline method. To this end, we randomly selected four images from the COCO val2017 dataset. The pose estimation results from the baseline method are displayed in the top row. In the bottom row, we show the refined keypoint estimations produced by our SCAI method. As can be observed, our method successfully corrects the keypoints circled in red, such as the left knee of the person in the third column. This indicates the effectiveness of our approach in refining pose estimation results.




Figure 6. Four examples of refinement of predicted keypoints. The top row is the original estimation. The bottom row is the estimation from SCAI.

In our self-correctable network design, we aim to correct the prediction output so as to minimize the self-referential error. Figure 7 shows two examples where the self-referential error can be used as a feedback reference to guide the correction network to pull the prediction result toward the point with the minimum self-referential error. More importantly, this point is often very close to the ground truth (circle) since the self-referential feedback error is highly correlated with the keypoint prediction error.

5. Further Discussions and Summary of Unique Contributions

In this section, we provide further discussion of our proposed SCAI method compared with related work and sum-




Figure 7. Two examples of corrected keypoints from SCAI. The distribution shows keypoints error by the ground truth (red dot), prediction network (yellow square) and our method SCAI (green triangle), where the blue curves represent errors of randomly selected from correction process.

marize its unique contributions.

(1) Unique Differences from Related Work. This work is related to the SCIO [15]. In our design, we borrowed the structural grouping of the human body keypoints from the paper. We also used its verification network to pre-train our fitness feedback network. Compared to this paper, the main part of our algorithm, namely, the self-correctable and adaptable inference method, including the correction network, fitness feedback network, and inference-time network optimization, are totally new. From the experimental results, we can see that our SCAI algorithm outperforms the method in Kan *et al.* [15] by large margins. Kan *et al.* [15] performed a local search to refine the prediction result with very high computational complexity. In this work, we learned a correction network to correct the prediction error which is much more effective and has much lower computational complexity. Also, as mentioned in section 4.3, our method performs more accurate keypoint estimation on the optimization process during the inference stage.

This work is also related to cycle consistency [44], reciprocal learning [29] and dual learning [11, 38]. Compared to these methods, our method is uniquely different and novel since our method establishes the correctable and adaptable inference with the ability to correct the prediction result and update the network model during the inference stage. However, such correction processes and adaptable inference are not available in the above existing methods and they only used cycle constraints for model training and testing.

(2) Algorithm complexity. Table 6 demonstrates that our proposed SCAI method introduces additional computational complexity. Specifically, the FFN and correction network increase the complexity of the baseline pose estimation model. It is worth noting that these two networks operate on the previously predicted heatmaps and thus their net-

work structures and complexity are relatively modest. The self-referential adaptation requires updating the correction network multiple times for the entire batch, thereby increasing the overall computational burden. However, compared to SCIO, our model offers improved inference speed. In our future work, we aim to explore more efficient methods of self-adaptable inference.

Table 6. Complexity analysis on COCO val set.

Method	Parameters (M)	Speed (fps)
HRNet [30]	64	125
SCIO [15]	193	72
SCAI (Ours)	357	81

(3) Summary of contributions. The major contributions of this work can be summarized as follows: (a) We demonstrate that it is theoretically possible to learn a feedback-correction network to refine the prediction results of a well-trained network, outperforming the SOTA and SCIO by 1.4%, which is quite significant. (b) We have introduced a correction network which is able to correct the prediction error for the test sample guided by the self-referential feedback error. This error was generated by a learned fitness feedback network. We found that this self-referential error is highly correlated with the actual network prediction error. (c) Using the self-referential error, we have introduced a new loss function to perform quick adaptation and optimization of the correction network during the inference stage. (d) We apply the proposed self-correctable and adaptable inference method to human pose estimation and have achieved remarkable performance gain and significant improvement of generalization capability of the pose estimation network.

6. Conclusion

In this work, we have developed a self-correctable and adaptable inference method to address the generalization challenge of network prediction and use human pose estimation as an example to demonstrate its effectiveness and performance. We have introduced the self-referential feedback error on the network input samples by constructing a feedback fitness network, so that it is able to evaluate if the prediction is accurate or not, without the need to know the ground truth. Guided by the self-referential error, we learn a prediction correction network which is able to adjust the prediction result during the dynamic inference process. The self-referential error is employed as a loss function for network adaptation during the inference phase. Our comprehensive experimental results on human pose estimation attest to the ability of the SCAI method to enhance both the generalization ability and performance of human pose estimation to a significant degree.

References

- [1] Timur M. Bagautdinov, Alexandre Alahi, François Fleuret, Pascal Fua, and Silvio Savarese. Social scene understanding: End-to-end multi-person action localization and collective activity recognition. In *CVPR*, pages 3425–3434, 2017. **1**
- [2] Zhe Cao, Tomas Simon, Shih-En Wei, and Yaser Sheikh. Realtime multi-person 2d pose estimation using part affinity fields. In *CVPR*, pages 1302–1310, 2017. **1**
- [3] Joao Carreira, Pulkit Agrawal, Katerina Fragkiadaki, and Jitendra Malik. Human pose estimation with iterative error feedback. In *CVPR*, pages 4733–4742, 2016. **2**
- [4] Yilun Chen, Zhicheng Wang, Yuxiang Peng, Zhiqiang Zhang, Gang Yu, and Jian Sun. Cascaded pyramid network for multi-person pose estimation. In *CVPR*, pages 7103–7112, 2018. **1, 2, 6**
- [5] Bowen Cheng, Bin Xiao, Jingdong Wang, Honghui Shi, Thomas S. Huang, and Lei Zhang. Higherhrnet: Scale-aware representation learning for bottom-up human pose estimation. In *CVPR*, pages 5385–5394, 2020. **6**
- [6] Guilhem Chéron, Ivan Laptev, and Cordelia Schmid. P-cnn: Pose-based cnn features for action recognition. In *ICCV*, pages 3218–3226, 2015. **1**
- [7] Ahmed Elhayek, Edilson de Aguiar, Arjun Jain, Jonathan Tompson, Leonid Pishchulin, Mykhaylo Andriluka, Christoph Bregler, Bernt Schiele, and Christian Theobalt. Efficient convnet-based marker-less motion capture in general scenes with a low number of cameras. In *CVPR*, pages 3810–3818, 2015. **1**
- [8] Hao-Shu Fang, Shuqin Xie, Yu-Wing Tai, and Cewu Lu. Rmpe: Regional multi-person pose estimation. In *ICCV*, pages 2334–2343, 2017. **6**
- [9] Mihai Fieraru, Anna Khoreva, Leonid Pishchulin, and Bernt Schiele. Learning to refine human pose estimation. In *CVPRW*, pages 205–214, 2018. **2**
- [10] Thomas Golda, Tobias Kalb, Arne Schumann, and Jürgen Beyerer. Human pose estimation for real-world crowded scenarios. In *AVSS*, pages 1–8. IEEE, 2019. **6**
- [11] Di He, Yingce Xia, Tao Qin, Liwei Wang, Nenghai Yu, Tie-Yan Liu, and Wei-Ying Ma. Dual learning for machine translation. In *NeurIPS*, pages 820–828, 2016. **8**
- [12] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *ICCV*, pages 2961–2969, 2017. **1, 6**
- [13] 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 *CVPR*, pages 5699–5708, 2020. **1, 2, 6**
- [14] Shaoli Huang, Mingming Gong, and Dacheng Tao. A coarse-fine network for keypoint localization. In *ICCV*, pages 3047–3056, 2017. **6**
- [15] Zhehan Kan, Shuoshuo Chen, Zeng Li, and Zhihai He. Self-constrained inference optimization on structural groups for human pose estimation. In *ECCV*, volume 13665, pages 729–745. Springer, 2022. **2, 4, 6, 8**
- [16] Rawal Khirodkar, Visesh Chari, Amit Agrawal, and Amrith Tyagi. Multi-instance pose networks: Rethinking top-down pose estimation. In *ICCV*, pages 3122–3131, October 2021. **6**
- [17] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. **6**
- [18] Dangwei Li, Xiaotang Chen, Zhang Zhang, and Kaiqi Huang. Pose guided deep model for pedestrian attribute recognition in surveillance scenarios. In *ICME*, pages 1–6. IEEE, 2018. **1**
- [19] Jia Li, Wen Su, and Zengfu Wang. Simple pose: Rethinking and improving a bottom-up approach for multi-person pose estimation. In *AAAI*, volume 34, pages 11354–11361, 2020. **2**
- [20] Jiefeng Li, Can Wang, Hao Zhu, Yihuan Mao, Hao-Shu Fang, and Cewu Lu. Crowdpose: Efficient crowded scenes pose estimation and a new benchmark. In *CVPR*, pages 10863–10872, 2019. **2, 5, 6**
- [21] 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. **6**
- [22] Yizhuo Li, Miao Hao, Zonglin Di, Nitesh B. Gundavarapu, and Xiaolong Wang. Test-time personalization with a transformer for human pose estimation. In *NeurIPS*, pages 2583–2597, 2021. **2**
- [23] 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 *ECCV*, pages 740–755. Springer, 2014. **5**
- [24] Zhengxiong Luo, Zhicheng Wang, Yan Huang, Liang Wang, Tieniu Tan, and Erjin Zhou. Rethinking the heatmap regression for bottom-up human pose estimation. In *CVPR*, pages 13264–13273, 2021. **2**
- [25] Gyeongsik Moon, Ju Yong Chang, and Kyoung Mu Lee. Posefix: Model-agnostic general human pose refinement network. In *CVPR*, pages 7773–7781, 2019. **2, 6**
- [26] 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 *CVPR*, pages 3711–3719, 2017. **1**
- [27] Helge Rhodin, Victor Constantin, Isinsu Katircioglu, Mathieu Salzmann, and Pascal Fua. Neural scene decomposition for multi-person motion capture. In *CVPR*, pages 7703–7713, 2019. **1**
- [28] Kai Su, Dongdong Yu, Zhenqi Xu, Xin Geng, and Changhu Wang. Multi-person pose estimation with enhanced channel-wise and spatial information. In *CVPR*, pages 5674–5682. Computer Vision Foundation / IEEE, 2019. **1, 2, 6**
- [29] Hao Sun, Zhiqun Zhao, and Zhihai He. Reciprocal learning networks for human trajectory prediction. In *CVPR*, pages 7414–7423, 2020. **8**
- [30] Ke Sun, Bin Xiao, Dong Liu, and Jingdong Wang. Deep high-resolution representation learning for human pose estimation. In *CVPR*, pages 5693–5703, 2019. **1, 4, 5, 6, 8**
- [31] Xiao Sun, Bin Xiao, Fangyin Wei, Shuang Liang, and Yichen Wei. Integral human pose regression. In *ECCV*, pages 536–553, 2018. **2**

- [32] Yu Sun, Xiaolong Wang, Zhuang Liu, John Miller, Alexei A. Efros, and Moritz Hardt. Test-time training with self-supervision for generalization under distribution shifts. In *ICML*, volume 119, pages 9229–9248. PMLR, 2020. [2](#)
- [33] Hsiao-Yu Tung, Hsiao-Wei Tung, Ersin Yumer, and Katerina Fragkiadaki. Self-supervised learning of motion capture. In *NeurIPS*, pages 5236–5246, 2017. [2](#)
- [34] Dequan Wang, Evan Shelhamer, Shaoteng Liu, Bruno A. Olshausen, and Trevor Darrell. Tent: Fully test-time adaptation by entropy minimization. In *ICLR*. OpenReview.net, 2021. [2](#)
- [35] Jian Wang, Xiang Long, Yuan Gao, Errui Ding, and Shilei Wen. Graph-pcnn: Two stage human pose estimation with graph pose refinement. In *ECCV*, pages 492–508. Springer, 2020. [2](#), [6](#)
- [36] Manchen Wang, Joseph Tighe, and Davide Modolo. Combining detection and tracking for human pose estimation in videos. In *CVPR*, pages 11085–11093, 2020. [1](#)
- [37] Jianchao Wu, Limin Wang, Li Wang, Jie Guo, and Gangshan Wu. Learning actor relation graphs for group activity recognition. In *CVPR*, pages 9964–9974, 2019. [1](#)
- [38] Yingce Xia, Tao Qin, Wei Chen, Jiang Bian, Nenghai Yu, and Tie-Yan Liu. Dual supervised learning. In *ICML*, volume 70 of *Proceedings of Machine Learning Research*, pages 3789–3798. PMLR, 2017. [8](#)
- [39] Bin Xiao, Haiping Wu, and Yichen Wei. Simple baselines for human pose estimation and tracking. In *ECCV*, pages 466–481, 2018. [1](#), [4](#), [5](#), [6](#)
- [40] Nan Xue, Tianfu Wu, Gui-Song Xia, and Liangpei Zhang. Learning local-global contextual adaptation for multi-person pose estimation. In *CVPR*, pages 13065–13074, 2022. [2](#)
- [41] Yiding Yang, Zhou Ren, Haoxiang Li, Chunlun Zhou, Xinchao Wang, and Gang Hua. Learning dynamics via graph neural networks for human pose estimation and tracking. In *CVPR*, pages 8074–8084, 2021. [1](#)
- [42] 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. [6](#)
- [43] Feng Zhang, Xiatian Zhu, Hanbin Dai, Mao Ye, and Ce Zhu. Distribution-aware coordinate representation for human pose estimation. In *CVPR*, pages 7091–7100, 2020. [6](#)
- [44] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A. Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *ICCV*, pages 2242–2251, 2017. [8](#)