

Dite-HRNet: Dynamic Lightweight High-Resolution Network for Human Pose Estimation

Qun Li¹, Ziyi Zhang¹, Fu Xiao^{1*}, Feng Zhang¹ and Bir Bhanu²

¹School of Computer Science, Nanjing University of Posts and Telecommunications, Nanjing, China

²Department of Electrical and Computer Engineering, University of California at Riverside, CA, USA
{liqun, xiaof, zhangfeng01}@njupt.edu.cn, zhangziyi_njupt@hotmail.com, bhanu@ee.ucr.edu

Abstract

A high-resolution network exhibits remarkable capability in extracting multi-scale features for human pose estimation, but fails to capture long-range interactions between joints and has high computational complexity. To address these problems, we present a Dynamic lightweight High-Resolution Network (Dite-HRNet), which can efficiently extract multi-scale contextual information and model long-range spatial dependency for human pose estimation. Specifically, we propose two methods, dynamic split convolution and adaptive context modeling, and embed them into two novel lightweight blocks, which are named dynamic multi-scale context block and dynamic global context block. These two blocks, as the basic component units of our Dite-HRNet, are specially designed for the high-resolution networks to make full use of the parallel multi-resolution architecture. Experimental results show that the proposed network achieves superior performance on both COCO and MPII human pose estimation datasets, surpassing the state-of-the-art lightweight networks. Code is available at: <https://github.com/ZiyiZhang27/Dite-HRNet>.

1 Introduction

Human pose estimation is a task that requires both accuracy and efficiency, especially when it encounters real-time applications on devices with limited resources. In human pose estimation, high-resolution representation is necessary for accuracy, which brings difficulties to achieve high efficiency.

While High-Resolution Network (HRNet) [Sun *et al.*, 2019] achieves remarkable performance in human pose estimation, it has high computational complexity. To make it lightweight, Small HRNet [Wang *et al.*, 2021] reduces the width and depth of the network, but such reduction introduces performance degradation. Lightweight High-Resolution Network (Lite-HRNet) [Yu *et al.*, 2021] shows that lightweight high-resolution network can achieve satisfactory performance by replacing each residual block in Small HRNet with an efficient CNN block. However, it remains unclear whether

such static block that is independent of the inputs can provide the optimal trade-off between the network performance and complexity for high-resolution networks, as some operations may have different effects at different locations of the network and with inputs of different sizes [Cui *et al.*, 2021]. The problem is that input-dependent component units may perform more efficiently in high-resolution networks than the input-independent counterparts. Our work focuses on designing dynamic blocks to develop a lightweight high-resolution network with dynamic representations.

Recent research [Wang *et al.*, 2018] has demonstrated the benefits of long-range spatial dependency in human pose estimation. The general concept of long-range spatial dependency is the global understanding of spatial information in a large field of view. Previous high-resolution networks [Sun *et al.*, 2019; Cheng *et al.*, 2020; Wang *et al.*, 2021; Yu *et al.*, 2021] mainly rely on the deeply stacked convolution layers in parallel branches that extract multi-scale features to build spatial dependency, but the capacity of lightweight networks can be severely affected by the restricted widths and depths. Hence, another problem is how to enhance lightweight high-resolution networks with more efficient methods for modeling long-range spatial dependency.

To address the above problems, we present a Dynamic lightweight High-Resolution Network (Dite-HRNet). We adopt the architecture similar to HRNet [Sun *et al.*, 2019] and develop two dynamic lightweight blocks to improve the overall efficiency. First, we propose a Dynamic Split Convolution (DSC), which can dynamically extract multi-scale contextual information and optimize the trade-off between its capacity and complexity by two hyper-parameters, and thus is more effective and flexible than the standard convolution. Then, we introduce long-range spatial dependency into the high-resolution network by designing a novel Adaptive Context Modeling (ACM) method that enables the model to learn both local and global patterns of human poses. Finally, we embed DSC and ACM into two dynamic lightweight blocks, which are specially designed for the high-resolution network to make full use of the parallel multi-resolution architecture, as the basic component units of our Dite-HRNet.

The main contributions of this paper can be summarized as follows: (1) We propose a novel lightweight network named Dite-HRNet that can efficiently extract multi-scale contextual information and model long-range spatial dependency for hu-

*Contact Author

man pose estimation. (2) We design a DSC method to build multi-scale contextual relationship and an ACM method to further strengthen long-range spatial dependency. In addition, we accomplish a joint embedding of DSC and ACM into two dynamic lightweight blocks, which are designed as the basic components of Dite-HRNet. (3) In our experiments, Dite-HRNet achieves superior trade-off between the network performance and complexity on both COCO and MPII human pose estimation datasets.

2 Related Work

Lightweight Human Pose Estimation. Recent studies [Wang *et al.*, 2021; Yu *et al.*, 2021] focus on improving the efficiency of human pose estimation networks. Small HRNet [Wang *et al.*, 2021] reduces the width and depth of the original large version of HRNet [Sun *et al.*, 2019]. Lite-HRNet [Yu *et al.*, 2021] exploits the potential of lightweight high-resolution network by replacing each residual block in the Small HRNet with a conditional channel weighting block, which adopts channel attention mechanism [Hu *et al.*, 2020] by replacing high-cost 1×1 convolutions in shuffle block [Ma *et al.*, 2018] with channel weighting operations.

Efficient CNN Blocks. Efficient CNN blocks have been widely used in many efficient CNN architectures [Howard *et al.*, 2017; Howard *et al.*, 2019; Zhang *et al.*, 2018], which aim to maximize the capacity of models under limited computational cost. MobileNetV2 [Sandler *et al.*, 2018] introduces the inverted residual block, which improves both accuracy and efficiency over traditional bottleneck blocks. Shuffle block in ShuffleNetV2 [Ma *et al.*, 2018] performs convolutions only on half of the channels with a competitive performance, due to the channel split operation that splits the features by channels and the channel shuffle operation that enhances the information exchange across channels. MixNet [Tan and Le, 2019] combines multiple convolution layer into a mixed depth-wise convolution, which is a drop-in replacement of vanilla depth-wise convolution in CNN blocks. Sandglass block [Zhou *et al.*, 2020] flips the structure of the inverted residual block, reducing the information loss without any additional computational cost.

Spatial Dependency Modeling. Long-range spatial dependency can be spontaneously captured in the repeated convolution layers, which have better performance with large convolution kernels. However, deeply stacked convolution layers and large convolution kernels are both costly and inflexible to be integrated in lightweight networks. Non-local network [Wang *et al.*, 2018] adopts a self-attention mechanism to model pixel-wise spatial relations in a single layer, which has better flexibility but still high complexity. Global context network [Cao *et al.*, 2019] simplifies the non-local network to a lightweight structure with almost no performance loss.

Dynamic CNN Architectures. There are two main stream methods for dynamic CNN architecture design. One is dynamic network structure, such as PP-LCNet [Cui *et al.*, 2021], which dynamically conducts different operations based on the inputs of different sizes in different positions of the network. And the other is dynamic convolution kernel, such as dynamic convolution [Chen *et al.*, 2020] and CondConv [Yang

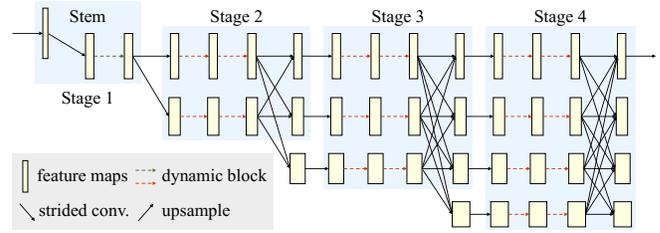


Figure 1: Overall architecture of Dite-HRNet. It uses HRNet [Sun *et al.*, 2019] as backbone, adapted by two proposed dynamic blocks. The cross-resolution modules are marked with blue areas.

Stage	Operator type	#Modules		#Branches
		Dite-HRNet-18	Dite-HRNet-30	
1 st	3×3 strided conv ($\times 1$) DGC block ($\times 1$)	1	1	1
2 nd	DMC block ($\times 2$) multi-scale fusion ($\times 1$)	2	3	2
3 rd	DMC block ($\times 2$) multi-scale fusion ($\times 1$)	4	8	3
4 th	DMC block ($\times 2$) multi-scale fusion ($\times 1$)	2	3	4

Table 1: Structure of Dite-HRNet. DGC = Dynamic Global Context. DMC = Dynamic Multi-scale Context. # = number of. \times denotes repetitions of each operator in one cross-resolution module.

et al., 2019], which mix different convolution kernels by generating weights through the attention mechanism. In order to exploit more efficient representations, our proposed Dite-HRNet embodies its dynamic state not only in the network structure but also in the convolution kernels.

3 Dite-HRNet

As shown in Figure 1, Dite-HRNet is a 4-stage network, consisting of one high-resolution main branch with the highest resolution and three high-to-low resolution branches that are added to the network one by one in parallel at the beginning of each new stage. Each newly added branch has half the resolution and twice the number of channels compared to the previously added branch. As shown in Table 1, among all four stages of Dite-HRNet, the first stage, also regarded as the stem, contains a 3×3 strided convolution and a Dynamic Global Context (DGC) block on the main branch. Each subsequent stage consists of a series of cross-resolution modules which are composed of two Dynamic Multi-scale Context (DMC) blocks and a multi-scale fusion layer that exchanges information across all branches. The main branch with the highest resolution maintains a high-resolution representation, which provides the final output of the backbone network for the subsequent pose estimation. For fair comparisons, we presents two instantiations of our network, Dite-HRNet-18 and Dite-HRNet-30, which have the network widths and depths corresponding to Lite-HRNet-18 [Yu *et al.*, 2021] and Lite-HRNet-30 [Yu *et al.*, 2021], respectively.

Block Design. As shown in Figure 2, our DMC block and DGC block share similar overall structure, applying the channel split, feature concatenation and channel shuffle operation from ShuffleNetV2 [Ma *et al.*, 2018] to aggregate different

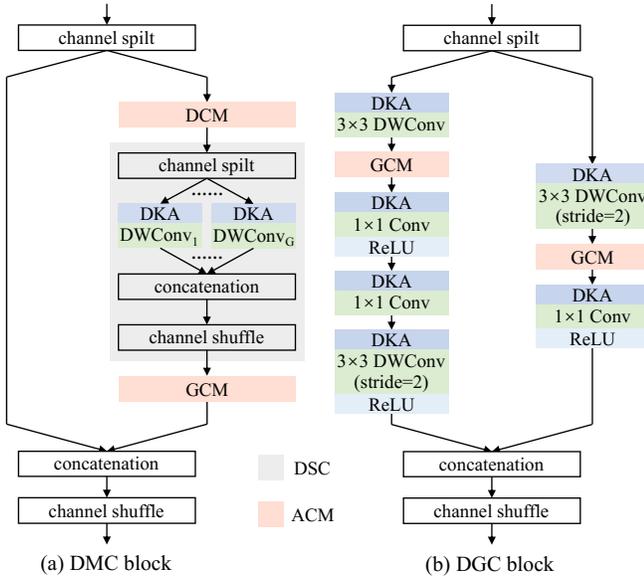


Figure 2: Structures of Dynamic Multi-scale Context (DMC) block and Dynamic Global Context (DGC) block. Dynamic Kernel Aggregation (DKA) is applied to Dynamic Split Convolution (DSC) in the DMC block and each convolution in the DGC block. DCM (Dense Context Modeling) and GCM (Global Context Modeling) are two instantiations of our proposed Adaptive Context Modeling (ACM).

features extracted by different layers. One difference between the two blocks is that the DMC block applies a sequence of layers on half of the channels, while the DGC block applies two different sequences of layers on all two groups of channels. The sequence of layers in the DMC block contains one Dense Context Modeling (DCM) operation, one DSC and one Global Context Modeling (GCM). Both the DCM and GCM are the instantiations of the ACM method. In the DGC block, one 3×3 strided depth-wise convolution, one GCM and one 1×1 convolution are performed on one group of channels, while one 3×3 depth-wise convolution, one GCM, one 1×1 convolution and one 3×3 strided depth-wise convolution are performed on the other group of channels. Each convolution in the DGC block and the DSC layer generates the convolution kernel via Dynamic Kernel Aggregation (DKA).

Configuration. In order to demonstrate the superior efficiency of our method, we optimize our networks, which are characterized by two hyper-parameters in DSC, to have similar model sizes and computational cost like in Lite-HRNet [Yu *et al.*, 2021]. We configure the hyper-parameters separately on different resolution branches. Our default configuration choice is shown in Table 5 and used for the subsequent comparison experiments.

3.1 Dynamic Split Convolution (DSC)

Split-Concat-Shuffle (SCS) Module. Large convolution kernels can bring broad receptive fields, thereby enhancing the long-range spatial dependency. However, using large kernel sizes on all filters in the network not only leads to a high computational cost, but also fails to capture the local pixel-wise information. Therefore, we introduce the SCS module

to extract contextual information by multiple kernels of different sizes and integrate them in a single convolution layer.

In the SCS module, channels are first split into multiple groups equally, and depth-wise convolutions with different kernel sizes are applied to each group of channels in parallel. The output of the convolution on each group is formally defined as follows:

$$Y_i = DWConv(K_i \times K_i | C/G)(X_i), \quad (1)$$

$$K_i = 2i + 1, \quad i \in [1, G],$$

where X_i and Y_i denote the input and output of the depth-wise convolution on the i^{th} group of channels, respectively. $DWConv(K_i \times K_i | C/G)(\cdot)$ is the depth-wise convolution with the kernel size $K_i \times K_i$ and channel dimension C/G , where C denotes the total number of channels among the groups and G denotes the number of groups.

After the depth-wise convolutions, the grouped features are concatenated together. To further integrate the separated information at different scales, the channel shuffle operation [Ma *et al.*, 2018] is used at the bottom of the SCS module.

Note that the SCS module does not expand the width of the network, as it only splits channels into different groups and performs different convolution operations on them in parallel. The computational efficiency is acceptable and can be optimized by the hyper-parameter G .

Dynamic Kernel Aggregation (DKA). To make the SCS module learn rich contextual information even with small convolution kernels, we introduce a DKA operation which strengthens the input-dependency of convolution kernels by dynamically aggregating multiple kernels via kernel-wise attention weights based on the input images.

Standard convolution kernel is defined by a weight matrix w with 4 dimensions which decide the kernel size and input/output channels, respectively. Instead of concatenating the output features of different convolutions, we aggregate the kernel weight matrices $\{w_i\}$ before calculating the convolution results, to dynamically generate different convolution kernels for different inputs. The DKA operation computes attention weights over different convolution kernels, and then applies the element-wise product to the attention weights and the kernel weights. We define the DKA operation as follows:

$$Y = W^T(X)X, \quad (2)$$

$$W(X) = \sum_i^N a_i(X)w_i,$$

where $a_i(X)$ is the attention weight for the i^{th} convolution kernel, and $W(X)$ is the aggregated weight matrix of N convolution kernels. The input-dependent attention weights $a_i(X)$ are computed from the input X as follows:

$$a_i(X) = Sigmoid(FC(ReLU(FC(GAP(X))))), \quad (3)$$

where $GAP(\cdot)$ represents the global average pooling, and $FC(\cdot)$ represents the full-connected layer. Two functions $Sigmoid(\cdot)$ and $ReLU(\cdot)$ are used right after two full-connected layers for non-linear activation.

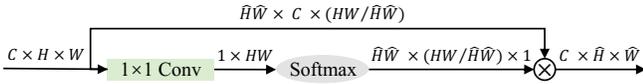


Figure 3: Procedure of adaptive context pooling. C denotes the number of channels. H and W denote the current height and width of features. \hat{H} and \hat{W} denote the target height and width of features. \otimes denotes matrix multiplication.

Since the DKA operation takes place before the calculation of convolution results, the aggregated kernel performs only one convolution operation on each input feature map without expanding the network width. For a standard convolution with kernel size $K \times K$ and input size $C \times H \times W$, DKA of N kernels only introduces minor additional computational cost $HWC + C^2/4 + CN/4$ from the standard convolution cost HWC^2K^2 . The number of aggregated kernels N is another hyper-parameter for the efficiency optimization.

Complementation. The high computational efficiency makes the DKA operation an excellent complementation to the SCS module. Therefore, we combine the above two methods into one single convolution, which is named DSC. With two hyper-parameters G and N from the SCS module and the DKA operation, respectively, it is relatively easier for DSC to optimize the trade-off between the capacity and complexity of the convolution layers.

3.2 Adaptive Context Modeling (ACM)

The ACM method can be abstracted into following three steps: (a) adaptive context pooling, as shown in Figure 3, which creates a shortcut connection with a 1×1 channel pooling convolution and a softmax layer to form a context mask, and then applies the mask on the feature maps through a series of resolution-adaptive transformations to obtain spatial contextual features; (b) context shifting, which realigns the spatially related contextual features together by two 1×1 convolutions with non-linear activation; (c) context weighting, which adopts the element-wise weighting operation between input features and shifted context features to model the contextual relationships of the corresponding features. This abstraction of the ACM method can be defined as:

$$Y = Weight(X, Shift(ACPool(H, W)(X))), \quad (4)$$

where $ACPool(H, W)(X)$ denotes the adaptive context pooling which pools the input features X to a certain output size $H \times W$, $Shift(\cdot)$ denotes the context shifting, and $Weight(\cdot)$ denotes the context weighting.

For the high-resolution network, we propose two instantiations of the ACM method, DCM and GCM, which take the advantage of the parallel multi-resolution architecture.

Dense Context Modeling (DCM). We introduce a DCM operation to densely model the spatial contextual relationships of the features from all resolution branches of one stage. At the n^{th} stage, the input features from all n branches are pooled to the lowest resolution $H_n \times W_n$. Then, all pooled features are concatenated together, so that the context shifting can be densely performed on the parallel contextual features. In the end, the shifted contextual features are upsampled to

the corresponding resolutions and distributed back to the corresponding branches for the subsequent context weighting. This instantiation implements the ACM as:

$$\begin{cases} \hat{X}_k = ACPool(H_n, W_n)(X_k), \\ \tilde{X} = Shift(Cat([\hat{X}_1, \dots, \hat{X}_{n-1}, X_n])), \\ Y_k = \begin{cases} Weight(X_k, Upsamp(\tilde{X}_k)), & 1 \leq k \leq n-1, \\ Weight(X_k, \tilde{X}_k), & k = n, \end{cases} \end{cases} \quad (5)$$

where $Cat(\cdot)$ and $Upsamp(\cdot)$ denote the feature concatenation and upsampling, respectively. X_k represents the input tensor with the k^{th} highest resolution. \hat{X} represents the pooled tensor from the k^{th} branch. \tilde{X}_k represents the shifted tensor, which is distributed to the k^{th} branch as \tilde{X}_k . Y_k represents the corresponding k^{th} output tensor.

Global Context Modeling (GCM). To separately model the global spatial dependency at each resolution, we apply a GCM operation on each branch of the network. It is a instantiation of the ACM when the adaptive context pooling has the output size 1×1 . The output features of the GCM operation on the k^{th} branch is defined as follows:

$$Y_k = Weight(X_k, Shift(ACPool(1, 1)(X_k))), \quad (6)$$

$$1 \leq k \leq n.$$

The GCM operation captures the spatial relationship of all features with the same resolution in the global aspect containing abundant contextual information, while the DCM operation captures the spatial relationships of all features with different resolutions in a moderate aspect containing more pixel-wise information. Meanwhile, both two operations increase the information exchange across features, so that can be better substitutes for the 1×1 convolutions in the shuffle block [Ma *et al.*, 2018] than the channel weighting operations in Lite-HRNet [Yu *et al.*, 2021].

4 Experiments

4.1 Implementation Details

Datasets and Evaluation Metrics. The COCO dataset [Lin *et al.*, 2014] has images over 200K and 250K person instances, each with a label of 17 keypoints. We train our networks on the train2017 set (contains 57K images and 150K person instances), and evaluate them on the val2017 set (contains 5K images) and test-dev2017 set (contains 20K images) by the Average Precision (AP) and Average Recall (AR) scores based on Object Keypoint Similarity (OKS). To further validate our networks, we also perform experiments on the MPII Human Pose dataset [Andriluka *et al.*, 2014], which contains about 25K images with 40K person instances, and evaluate the accuracy by the head-normalized Probability of Correct Keypoint (PCKh) score.

Training. The presented Dite-HRNet is trained on 8 GeForce RTX 3090 GPUs, with 32 samples per GPU. All parameters are updated by Adam optimizer with a base learning rate $2e^{-3}$. As for the data processing, we expand all human detection boxes to a fixed aspect ratio 4 : 3, and then crop the images with the detection boxes, which are resized to 256×192 or 384×288 for the COCO dataset, and 256×256

Method	Backbone	Pretrain	Input size	#Params (M)	GFLOPs	AP	AP ⁵⁰	AP ⁷⁵	AP ^M	AP ^L	AR
Large networks											
8-stage Hourglass [Newell <i>et al.</i> , 2016]	Hourglass	N	256 × 192	25.1	14.3	66.9	–	–	–	–	–
CPN [Chen <i>et al.</i> , 2018]	ResNet-50	Y	256 × 192	27.0	6.2	68.6	–	–	–	–	–
SimpleBaseline [Xiao <i>et al.</i> , 2018]	ResNet-50	Y	256 × 192	34.0	8.9	70.4	88.6	78.3	67.1	77.2	76.3
HRNet [Sun <i>et al.</i> , 2019]	HRNet-W32	N	256 × 192	28.5	7.1	73.4	89.5	80.7	70.2	80.1	78.9
UDP [Huang <i>et al.</i> , 2020]	HRNet-W32	Y	256 × 192	28.7	7.1	75.2	92.4	82.9	72.0	80.8	80.4
Small networks											
MobileNetV2 1× [Sandler <i>et al.</i> , 2018]	MobileNetV2	N	256 × 192	9.6	1.4	64.6	87.4	72.3	61.1	71.2	70.7
ShuffleNetV2 1× [Ma <i>et al.</i> , 2018]	ShuffleNetV2	N	256 × 192	7.6	1.2	59.9	85.4	66.3	56.6	66.2	66.4
Small HRNet [Wang <i>et al.</i> , 2021]	HRNet-W18	N	256 × 192	1.3	0.5	55.2	83.7	62.4	52.3	61.0	62.1
Lite-HRNet [Yu <i>et al.</i> , 2021]	Lite-HRNet-18	N	256 × 192	1.1	0.2	64.8	86.7	73.0	62.1	70.5	71.2
	Lite-HRNet-30	N	256 × 192	1.8	0.3	67.2	88.0	75.0	64.3	73.1	73.3
Dite-HRNet (Ours)	Dite-HRNet-18	N	256 × 192	1.1	0.2	<u>65.9</u>	<u>87.3</u>	<u>74.0</u>	<u>63.2</u>	<u>71.6</u>	<u>72.1</u>
	Dite-HRNet-30	N	256 × 192	1.8	0.3	68.3	88.2	76.2	65.5	74.1	74.2
MobileNetV2 1× [Sandler <i>et al.</i> , 2018]	MobileNetV2	N	384 × 288	9.6	3.3	67.3	87.9	74.3	62.8	74.7	72.9
ShuffleNetV2 1× [Ma <i>et al.</i> , 2018]	ShuffleNetV2	N	384 × 288	7.6	2.8	63.6	86.5	70.5	59.5	70.7	69.7
Small HRNet [Wang <i>et al.</i> , 2021]	HRNet-W18	N	384 × 288	1.3	1.2	56.0	83.8	63.0	52.4	62.6	62.6
Lite-HRNet [Yu <i>et al.</i> , 2021]	Lite-HRNet-18	N	384 × 288	1.1	0.4	67.6	87.8	75.0	64.5	73.7	73.7
	Lite-HRNet-30	N	384 × 288	1.8	0.7	70.4	88.7	77.7	67.5	76.3	76.2
Dite-HRNet (Ours)	Dite-HRNet-18	N	384 × 288	1.1	0.4	<u>69.0</u>	<u>88.0</u>	<u>76.0</u>	<u>65.5</u>	<u>75.5</u>	<u>75.0</u>
	Dite-HRNet-30	N	384 × 288	1.8	0.7	71.5	88.9	78.2	68.2	77.7	77.2

Table 2: Comparisons of results on the COCO val2017 set. Pretrain = pretrain the backbone on the ImageNet classification task. **Bold** indicates the best result and underline indicates the highest score with the lowest #Params or FLOPs.

Method	Backbone	Input size	#Params (M)	GFLOPs	AP	AP ⁵⁰	AP ⁷⁵	AP ^M	AP ^L	AR
Large networks										
SimpleBaseline [Xiao <i>et al.</i> , 2018]	ResNet-50	256 × 192	34.0	8.9	70.0	90.9	77.9	66.8	75.8	75.6
CPN [Chen <i>et al.</i> , 2018]	ResNet-Inception	384 × 288	–	–	72.1	91.4	80.0	68.7	77.2	78.5
HRNet [Sun <i>et al.</i> , 2019]	HRNet-W32	384 × 288	28.5	16.0	74.9	92.5	82.8	71.3	80.9	80.1
UDP [Huang <i>et al.</i> , 2020]	HRNet-W32	384 × 288	28.7	16.1	76.1	92.5	83.5	72.8	82.0	81.3
DARK [Zhang <i>et al.</i> , 2020]	HRNet-W48	384 × 288	63.6	32.9	76.2	92.5	83.6	72.5	82.4	81.1
Small networks										
MobileNetV2 1× [Sandler <i>et al.</i> , 2018]	MobileNetV2	384 × 288	9.8	3.3	66.8	90.0	74.0	62.6	73.3	72.3
ShuffleNetV2 1× [Ma <i>et al.</i> , 2018]	ShuffleNetV2	384 × 288	7.6	2.8	62.9	88.5	69.4	58.9	69.3	68.9
Small HRNet [Wang <i>et al.</i> , 2021]	HRNet-W18	384 × 288	1.3	1.2	55.2	85.8	61.4	51.7	61.2	61.5
Lite-HRNet [Yu <i>et al.</i> , 2021]	Lite-HRNet-18	384 × 288	1.1	0.4	66.9	89.4	74.4	64.0	72.2	72.6
	Lite-HRNet-30	384 × 288	1.8	0.7	69.7	90.7	77.5	66.9	75.0	75.4
Dite-HRNet (Ours)	Dite-HRNet-18	384 × 288	1.1	0.4	<u>68.4</u>	<u>89.9</u>	<u>75.8</u>	<u>65.2</u>	<u>73.8</u>	<u>74.4</u>
	Dite-HRNet-30	384 × 288	1.8	0.7	70.6	90.8	78.2	67.4	76.1	76.4

Table 3: Comparisons of results on the COCO test-dev2017 set. #Params and FLOPs are computed for pose estimation, and those for human detection are not included. **Bold** indicates the best result and underline indicates the highest score with the lowest #Params or FLOPs.

Method	#Params (M)	GFLOPs	PCKh
MobileNetV2 1× [Sandler <i>et al.</i> , 2018]	9.6	1.9	85.4
MobileNetV3 1× [Howard <i>et al.</i> , 2019]	8.7	1.8	84.3
ShuffleNetV2 1× [Ma <i>et al.</i> , 2018]	7.6	1.7	82.8
Small HRNet [Wang <i>et al.</i> , 2021]	<u>1.3</u>	0.7	80.2
Lite-HRNet-18 [Yu <i>et al.</i> , 2021]	1.1	0.2	86.1
Lite-HRNet-30 [Yu <i>et al.</i> , 2021]	1.8	<u>0.4</u>	<u>87.0</u>
Dite-HRNet-18 (Ours)	1.1	0.2	<u>87.0</u>
Dite-HRNet-30 (Ours)	1.8	<u>0.4</u>	87.6

Table 4: Comparisons of results on the MPII val set. **Bold** indicates the best result and underline indicates the second-best result.

for the MPII dataset. All images are used with data augmentations, including random rotations with factor 30, random scales with factor 0.25, and random flippings for both the COCO and MPII datasets. In addition, extra half body transformations are performed for the COCO dataset.

Testing. The two-stage top-down paradigm, which produces the person detection boxes and then predicts the person keypoints, is adopted for our testing process. The person boxes are predicted by the person detectors provided by SimpleBaseline [Xiao *et al.*, 2018] for the COCO dataset, while the standard testing strategy for the MPII dataset uses the provided person boxes. The heatmaps are estimated via 2D Gaussian, and then averaged for the original and flipped images. The highest heatmap locations in the heatmaps are adjusted by a quarter offset in the direction from the highest response to the second-highest response, to obtain the keypoint locations.

4.2 Results

Results on the COCO val2017 Set. As shown in Table 2, we compare the results of our Dite-HRNet-18 and Dite-HRNet-30 with other state-of-the-art pose estimation meth-

ods, including both large and small networks. Our networks are trained with two different input sizes 256×192 and 384×288 , achieving better trade-off between the accuracy and complexity than other methods. As compared to other small networks, our Dite-HRNet-30 achieves the highest AP score of 71.5 with the input size 384×288 . In particular, Dite-HRNet exceeds Small HRNet [Wang *et al.*, 2021] by 10 points of AP with only 40% GFLOPs. As compared to the large pose networks, our Dite-HRNet also achieves comparable or even higher accuracy with far smaller model sizes and lower computational complexity. Benefiting from the abundant contextual information, Dite-HRNet outperforms Lite-HRNet [Yu *et al.*, 2021] with the equivalent parameters and GFLOPs, while having the same network depths and input sizes.

Results on the COCO test-dev2017 Set. In Table 3, we report the results of comparisons of our networks and other state-of-the-art methods. Our Dite-HRNet achieves higher efficiency than both large and small networks. As compared to Lite-HRNet-18 [Yu *et al.*, 2021], Dite-HRNet-18 improves AP by 1.5 points and AR by 1.8 points with the equivalent model complexity. Our Dite-HRNet-30 achieves the highest AP of 70.6 among the small networks, outperforming Simple-Baseline [Xiao *et al.*, 2018] while having only 5% parameters and 8% GFLOPs. Despite the accuracy gap with other large networks, our networks are much more competitive in terms of the model size and computational complexity.

Results on the MPII val Set. The results of our networks compared with other lightweight networks are reported in Table 4. Our Dite-HRNet-18 improves PKCh@0.5 score by 0.9 points over Lite-HRNet-18 [Yu *et al.*, 2021] with the equivalent model complexity, has the same score but only half of GFLOPs compared to Lite-HRNet-30 [Yu *et al.*, 2021], and outperforms MobileNetV2 [Sandler *et al.*, 2018], MobileNetV3 [Howard *et al.*, 2019], ShuffleNetV2 [Ma *et al.*, 2018] and Small HRNet [Wang *et al.*, 2021] with much lower parameters and GFLOPs. With the same model size as Lite-HRNet-30, our Dite-HRNet-30 achieves the best result of 87.6 PKCh@0.5 among the lightweight networks.

Note that the accuracy improvements of Dite-HRNet-18 over Lite-HRNet-18 are more significant than that of Dite-HRNet-30 over Lite-HRNet-30 on all three sets. Moreover, our small version of Dite-HRNet achieves the proximate performance to the large version of Lite-HRNet. Therefore, our proposed methods are more effective for small networks, and also much more efficient than increasing the network depth.

4.3 Ablation Study

Hyper-parameters in DSC. To explore the best configuration of two hyper-parameters G , N in DSC, we conduct a 16-group ablation study on the COCO val2017 set. The results of our Dite-HRNet-18 with different configurations are shown in Table 5. On each branch of Dite-HRNet-18, parameters and GFLOPs increase as N , G increase. However, large N for the upper branches as well as large G for the lower branches have less impact on the model size and computational complexity, and also achieve more improvement on accuracy. It shows a better trade-off between the accuracy

		(a) #Params (M) / GFLOPs			
		N (Kernels)			
		1 1 1 1	4 4 2 1	1 2 4 4	4 4 4 4
G (Groups)	1 1 1 1	1.0 / 0.20	1.1 / 0.20	1.2 / 0.20	1.2 / 0.20
	1 1 2 4	1.1 / 0.20	1.1 / 0.20*	1.2 / 0.20	1.2 / 0.20
	4 2 1 1	1.1 / 0.23	1.2 / 0.24	1.2 / 0.24	1.3 / 0.24
	4 4 4 4	1.1 / 0.25	1.3 / 0.25	1.4 / 0.25	1.5 / 0.25

		(b) AP on the COCO val2017			
		N (Kernels)			
		1 1 1 1	4 4 2 1	1 2 4 4	4 4 4 4
G (Groups)	1 1 1 1	65.0	65.4	65.3	65.3
	1 1 2 4	65.3	65.9*	65.7	65.8
	4 2 1 1	65.5	65.4	65.3	65.7
	4 4 4 4	65.8	66.1	65.8	65.9

Table 5: Results of Dite-HRNet-18 with different configurations of two hyper-parameters in DSC, the number of groups G and the number of kernels N . The numbers in the header of each row or column denote the values of G or N on the highest to lowest resolution branches, respectively. The results of the model with the default configuration we choose are marked with *. The GFLOPs is computed with the input size 256×192 .

Method	#Params (M)	COCO		MPII	
		MFLOPs	AP	MFLOPs	PKCh
Lite-HRNet-18 [Yu <i>et al.</i> , 2021]	1.1	205.2	64.8	273.4	86.1
with ACM	1.0	210.6	65.1	280.6	86.5
with DSC	1.2	211.0	65.3	281.1	86.5
with ACM & DSC	1.1	215.5	65.6	287.2	86.7
Dite-HRNet-18 (Ours)	1.1	209.8	65.9	279.5	87.0

Table 6: Ablation studies on the COCO val2017 and MPII val sets. The MFLOPs is computed with the input size 256×192 for the COCO val2017 set and 256×256 for the MPII val set, respectively.

and complexity by setting G to 1, 1, 2, 4 and N to 4, 4, 2, 1 for the highest to lowest resolution branches, respectively.

Effectiveness of ACM and DSC. To further verify the effectiveness of our proposed methods, we perform ablation studies on the COCO val2017 and MPII val sets. We first use Lite-HRNet-18 [Yu *et al.*, 2021] as a baseline, to which our ACM and DSC are applied separately or together, and then compare the results with our Dite-HRNet-18. Table 6 shows that our ACM and DSC methods bring considerable accuracy improvements to Lite-HRNet-18 with negligible additional complexity for the overall model. Further, our lightweight blocks helps Dite-HRNet-18 achieve the best results and reduce the model complexity.

5 Conclusion

In this paper, in order to address the problems that previous high-resolution networks were input-independent and lacked long-range information, we presented a Dite-HRNet to dynamically extract feature representations for human pose estimation. Our network achieved impressive efficiency on both COCO and MPII human pose estimation datasets, due to the effectiveness of the DMC and DGC blocks, which performed input-dependent convolutions by embedding DSC and captured long-range information by embedding ACM. The proposed dynamic lightweight blocks can further expand to other networks that have multi-scale representations.

Acknowledgments

This work was partially supported by the National Science Fund for Distinguished Young Scholars of China under Grant No.62125203, the Key Program of the National Natural Science Foundation of China under Grant No.61932013, the National Natural Science Foundation of China under Grant No. 61906099 and No. 61906098.

References

- [Andriluka *et al.*, 2014] Mykhaylo Andriluka, Leonid Pishchulin, Peter Gehler, and Bernt Schiele. 2D human pose estimation: New benchmark and state of the art analysis. In *CVPR*, pages 3686–3693, 2014.
- [Cao *et al.*, 2019] Yue Cao, Jiarui Xu, Stephen Lin, Fangyun Wei, and Han Hu. GCNet: Non-local networks meet squeeze-excitation networks and beyond. In *ICCV*, pages 1971–1980, 2019.
- [Chen *et al.*, 2018] 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.
- [Chen *et al.*, 2020] Yinpeng Chen, Xiyang Dai, Mengchen Liu, Dongdong Chen, Lu Yuan, and Zicheng Liu. Dynamic convolution: Attention over convolution kernels. In *CVPR*, pages 11027–11036, 2020.
- [Cheng *et al.*, 2020] 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.
- [Cui *et al.*, 2021] Cheng Cui, Tingquan Gao, Shengyu Wei, Yuning Du, Ruoyu Guo, Shuilong Dong, Bin Lu, Ying Zhou, Xueying Lv, Qiwen Liu, Xiaoguang Hu, Dianhai Yu, and Yanjun Ma. PP-LCNet: A lightweight CPU convolutional neural network. arXiv:2109.15099, 2021.
- [Howard *et al.*, 2017] Andrew G. Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. MobileNets: Efficient convolutional neural networks for mobile vision applications. arXiv:1704.04861, 2017.
- [Howard *et al.*, 2019] Andrew Howard, Mark Sandler, Bo Chen, Weijun Wang, Liang-Chieh Chen, Mingxing Tan, Grace Chu, Vijay Vasudevan, Yukun Zhu, Ruoming Pang, Hartwig Adam, and Quoc Le. Searching for MobileNetV3. In *ICCV*, pages 1314–1324, 2019.
- [Hu *et al.*, 2020] Jie Hu, Li Shen, Samuel Albanie, Gang Sun, and Enhua Wu. Squeeze-and-Excitation networks. *IEEE Trans. Pattern Anal. Mach. Intell.*, 42(8):2011–2023, 2020.
- [Huang *et al.*, 2020] 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.
- [Lin *et al.*, 2014] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollar, and Lawrence Zitnick. Microsoft COCO: Common objects in context. In *ECCV*, pages 740–755, 2014.
- [Ma *et al.*, 2018] Ningning Ma, Xiangyu Zhang, Hai-Tao Zheng, and Jian Sun. ShuffleNet V2: Practical guidelines for efficient CNN architecture design. In *ECCV*, pages 122–138, 2018.
- [Newell *et al.*, 2016] Alejandro Newell, Kaiyu Yang, and Jia Deng. Stacked hourglass networks for human pose estimation. In *ECCV*, pages 483–499, 2016.
- [Sandler *et al.*, 2018] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. MobileNetV2: Inverted residuals and linear bottlenecks. In *CVPR*, pages 4510–4520, 2018.
- [Sun *et al.*, 2019] Ke Sun, Bin Xiao, Dong Liu, and Jingdong Wang. Deep high-resolution representation learning for human pose estimation. In *CVPR*, pages 5686–5696, 2019.
- [Tan and Le, 2019] Mingxing Tan and Quoc V. Le. Mix-Conv: Mixed depthwise convolutional kernels. In *BMVC*, pages 1–13, 2019.
- [Wang *et al.*, 2018] Xiaolong Wang, Ross B. Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *CVPR*, pages 7794–7803, 2018.
- [Wang *et al.*, 2021] Jingdong Wang, Ke Sun, Tianheng Cheng, Borui Jiang, Chaorui Deng, Yang Zhao, Dong Liu, Yadong Mu, Mingkui Tan, Xinggang Wang, Wenyu Liu, and Bin Xiao. Deep high-resolution representation learning for visual recognition. *IEEE Trans. Pattern Anal. Mach. Intell.*, 43(10):3349–3364, 2021.
- [Xiao *et al.*, 2018] Bin Xiao, Haiping Wu, and Yichen Wei. Simple baselines for human pose estimation and tracking. In *ECCV*, pages 472–487, 2018.
- [Yang *et al.*, 2019] Brandon Yang, Gabriel Bender, Quoc V. Le, and Jiquan Ngiam. CondConv: Conditionally parameterized convolutions for efficient inference. In *NeurIPS*, pages 1305–1316, 2019.
- [Yu *et al.*, 2021] Changqian Yu, Bin Xiao, Changxin Gao, Lu Yuan, Lei Zhang, Nong Sang, and Jingdong Wang. Lite-HRNet: A lightweight high-resolution network. In *CVPR*, pages 10435–10445, 2021.
- [Zhang *et al.*, 2018] Xiangyu Zhang, Xinyu Zhou, Mengxiao Lin, and Ran Sun. ShuffleNet: An extremely efficient convolutional neural network for mobile devices. In *CVPR*, pages 6848–6856, 2018.
- [Zhang *et al.*, 2020] 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.
- [Zhou *et al.*, 2020] Daquan Zhou, Qibin Hou, Yunpeng Chen, Jiashi Feng, and Shuicheng Yan. Rethinking bottleneck structure for efficient mobile network design. In *ECCV*, pages 680–697, 2020.