

Dynamic Group Convolution for Accelerating Convolutional Neural Networks

Zhuo Su^{1,*}, Linpu Fang^{2,*}, Wenxiong Kang²,
Dewen Hu³, Matti Pietikäinen¹, and Li Liu^{3,1,†}

¹ Center for Machine Vision and Signal Analysis, University of Oulu, Finland

² South China University of Technology, China

³ National University of Defense Technology, China

Abstract. Replacing normal convolutions with group convolutions can significantly increase the computational efficiency of modern deep convolutional networks, which has been widely adopted in compact network architecture designs. However, existing group convolutions undermine the original network structures by cutting off some connections permanently resulting in significant accuracy degradation. In this paper, we propose dynamic group convolution (DGC) that adaptively selects which part of input channels to be connected within each group for individual samples on the fly. Specifically, we equip each group with a small feature selector to automatically select the most important input channels conditioned on the input images. Multiple groups can adaptively capture abundant and complementary visual/semantic features for each input image. The DGC preserves the original network structure and has similar computational efficiency as the conventional group convolution simultaneously. Extensive experiments on multiple image classification benchmarks including CIFAR-10, CIFAR-100 and ImageNet demonstrate its superiority over the existing group convolution techniques and dynamic execution methods⁴. The code is available at <https://github.com/zhuogege1943/dgc>.

Keywords: Group convolution, dynamic execution, efficient network architecture

1 Introduction

Deep convolutional neural networks (CNNs) have achieved significant successes in a wide range of computer vision tasks including image classification [5], object detection [30], and semantic segmentation [36]. Earlier studies [20,11] found that deeper and wider networks could obtain better performance, which results in a large number of huge and complex models being designed in the community.

* Equal contributions. † Corresponding author: li.liu@oulu.fi

⁴ This work was partially supported by the Academy of Finland and the National Natural Science Foundation of China under Grant 61872379. The authors also wish to acknowledge CSC IT Center for Science, Finland, for computational resources.




Fig. 1. The average contributions from the input channels to several output channels at a certain layer in a trained DenseNet [20] are depicted. The X-axis and Y-axis indicate the input channels and their average contributions to the output channels. For a pair of input and output channel, an activation map is firstly obtained by convolution with the corresponding filter, and the average contribution from this input channel to the output channel is calculated as the average value of the activation map. It can be seen that such contributions vary a lot across different input-output pairs, which is the main observation behind the motivation of the proposed DGC.

However, these models are very compute-intensive making them impossible to be deployed on edge devices with strict latency requirements and limited computing resources. In recent years, more and more researchers turn to study network compression techniques or design computation-efficient architectures to solve this troublesome.

Group convolution, which was first introduced in AlexNet [25] for accelerating the training process across two GPUs, has been comprehensively applied in computation-efficient network architecture designs [55,37,52,4,45]. Standard group convolution equally split the input and output channels in a convolution layer into G mutually exclusive groups while performing normal convolution operation within individual groups, which reduces the computation burden by G times in theory. The predefined group partition in standard group convolution may be suboptimal, recent studies [48,19] further propose learnable group convolution that learns the connections between input and output features in each group during the training process.

By analyzing the existing group convolutions, it can be found that they have two key disadvantages: 1) They weaken the representation capability of the normal convolution by introducing sparse neuron connections and suffer from decreasing performance especially for those difficult samples; 2) They have fixed neuron connection routines, regardless of the specific properties of individual inputs. However, the dependencies among input and output channels are not fixed and vary with different input images, which can be observed in Fig. 1. Here, for two different input images, the average contributions from the input channels

to several output channels at a certain layer in a trained DenseNet [20] are depicted. Two interesting phenomenons can be found in Fig. 1. Firstly, an output channel may receive information from input channels with varying contributions depending on a certain image, some of them are negligible. Secondly, for a single input image, some input channels with such negligible contributions correspond to groups of output channels, *i.e.*, the corresponding connections could be cut off without influences on the final results. It indicates that it needs an adaptive selection mechanism to select which set of input channels to be connected with output channels in individual groups.

Motivated by the dynamical computation mechanism in dynamic networks [29,9,17], in this paper, we propose dynamic group convolution (DGC) to adaptively select the most related input channels for each group while keeping the full structure of the original networks. Specifically, we introduce a tiny auxiliary feature selector for each group to dynamically decide which part of input channels to be connected based on the activations of all of input channels. Multiple groups can capture different complementary visual/semantic features of input images, making the DGC powerful to learn plentiful feature representations. Note that the computation overhead added by the auxiliary feature selectors are negligible compared with the speed-up provided by the sparse group convolution operations. In addition, the proposed DGC is compatible with various exiting deep CNNs and can be easily optimized in an end-to-end fashion.

We embed the DGC into popular deep CNN models including ResNet [11], CondenseNet [19] and MobileNetV2 [44], and evaluate its effectiveness on three common image recognition benchmarks: CIFAR-10, CIFAR-100 and ImageNet. The experimental results indicate the DGC outperforms the exiting group convolution techniques and dynamic execution methods.

2 Related Work

Efficient Architecture Design. As special cases of sparsely connected convolution, group convolution and its extreme version, *i.e.* depth-wise separable convolution, are most popular modules employed in efficient architecture designs. AlexNet [25] firstly uses group convolution to handle the problem of memory limitation. ResNeXt [52] further applies group convolution to implement a set of transformations and demonstrates its effectiveness. A series of subsequent researches use group convolution or depth-wise separable convolution to design computation-efficient CNNs [15,44,55,37,54,45]. Instead of predefining the connection models, CondenseNet [19] and FLGC [48] propose to automatically learn the connections of group convolution during the training process. All these exiting group convolutions have fixed connections during inference, inevitably weakening the representation capability of the original normal convolutions due to the sparse structures. Our proposed DGC can effectively solve this troublesome by employing dynamic execution strategy that keeps sparse computation without undermining original network structures.




Fig. 2. Overview of a DGC layer. For a simple illustration, we set the number of channels as 8 for both input and output volume and the number of heads as 2. Each head is responsible for generating half of the output volume, with different part of the input volume, by respectively using a saliency guided channel selector (gate). Here, white blocks after the gate represents the corresponding channel would not take part in the convolution.

Network Compression. Generally, compression methods can be categorized into five types: quantization [3,31,7,43,1], knowledge distillation [14,23,41,53,32], low-rank decomposition [38,6,22,56], weight sparsification [10,26,51], and filter pruning [34,27,13,40]. Quantization methods accelerate deep CNNs by replacing high-precision float point operations with low-precision fixed point ones, which usually incurs significantly accuracy drop. Knowledge distillation methods aim to learn a small student model by mimicking the output or feature distributions of a larger teacher model. Low-rank decomposition methods reduce the computation by factorizing the convolution kernels. Weight sparsification methods removes individual connections between neural nodes resulting in irregular models, while filter pruning methods, which our method is most related to, directly cut off entire filters keeping the regular architectures. Generally, those compression methods usually need several rounds to obtain the final compact models. On the contrary, the proposed DGC can be easily optimized with any exiting networks in an end-to-end manner. Note that the filter pruning removes some channels after training, which loses the capabilities of the original CNNs permanently, while our method keeps the capabilities by the dynamic channel selection scheme.

Dynamic Network. In contrast to the one-size-fit-all paradigm, dynamic networks [49,28,50,33,42,18,8] dynamically execute different modules across a deep network for individual inputs. One kind of method [2,42,18] sets multiple classifiers at different places in a deep network and apply early exiting strategy to implement dynamic inference. Another line of work [49,28,47,29,50,33,8] learns auxiliary controllers or gating units to decide which parts of layers or channels could be skipped for extracting intermediate feature representations. Among those dynamic execution methods, our proposed DGC is mostly motivated by the channel or pixel-level ones [29,9,17]. Runtime Neural Pruning (RNP) [29] learns a RNN to generate binary policies and adaptively prune channels in convolutional layers. The RNN is trained by reinforcement learning with a policy

function, which is hard to converge and sensitive to hyper-parameters. Feature boosting and suppression (FBS) [9] generates decisions to dynamically prune a subset of output channels, which is prone to suffer from the *internal covariate shift* problem [21] due to its interference in batch normalization (BN), bringing instability to training. Channel gating neural network (CGNet) [17] uses a subset of input channels to generate a binary decision map, and skips unimportant computation in the rest of input channels. The CGNet actually works in a semi-dynamic way and uses a complex approximating non-differentiable gate function to learn the binary decisions. In contrast to these methods, our proposed DGC works better to achieve stable training with dynamic channel selection, and can be easily optimized using common SGD algorithms.

3 Dynamic Group Convolution

3.1 Group Convolution and Dynamic Channel Pruning Revisiting

Consider how a CNN layer works with groups, which can be generalized as:

$$\mathbf{x}' = [\hat{f}(\mathbf{x}^1), \hat{f}(\mathbf{x}^2), \dots, \hat{f}(\mathbf{x}^N)], \quad (1)$$

$$\hat{f}(\mathbf{x}^i) = f(\mathbf{x}^i, \boldsymbol{\theta}^i) \cdot \pi(\mathbf{x}^i, \boldsymbol{\phi}^i), \quad (2)$$

where $\mathbf{x} \in \mathbb{R}^{C \times H \times W}$, $\mathbf{x}' \in \mathbb{R}^{C' \times H' \times W'}$ are the input and output of the current layer respectively, $\mathbf{x}^i \subseteq \mathbf{x}$ is a subset of input channels that is divided into the *i*th group, N is the number of groups, f executes the conventional convolution using parameter $\boldsymbol{\theta}$, while the meaning of π is case-dependent, with its parameter $\boldsymbol{\phi}$, which would be discussed below.

In a standard convolution layer, $N = 1$ and $\pi \equiv \mathbf{1}$. While a standard group convolution (SGC) evenly allocates equal number of channels in \mathbf{x} to each group with hard assignments, thus reducing the computation cost by N times. On the other hand, studies show that although complemented with channel permutation before performing group convolution, SGC is not always optimal [19,48], and introduce a learning strategy to soften the hard assignments, leading to a better group allocation. Such soft assignments allow each channel in \mathbf{x} to automatically choose its fitted group(s). To achieve the soft assignments, common methods are to add a group-lasso loss to induce group-level weights sparsity or a learnable binary decision matrix [19,48]. However, these methods can not be easily transferred to dynamic versions due to their “staticness”.

In a convolution layer with dynamic execution [9,8,17,29] and $N = 1$, a on-the-fly gate score or saliency is pre-calculated through a prediction function (represented as π in (2)) for each output cell that needs to be computed by f . The cell could be an output channel [9,29] or a pixel in the output feature maps [8,17]. Computation of the output cell in f is skipped if the corresponding saliency is small enough. Denoting the saliency vector (output of π , with size equal to the number of output cells) as \mathbf{g} , sparsity for \mathbf{g} is implemented by reinforcement learning in [29], which is hard to be optimized. [8,9] associate \mathbf{g} with a lasso loss, which is in fact a L1 norm regularization:

$$L_{lasso} = \|\mathbf{g}_1, \mathbf{g}_2, \dots\|_1. \quad (3)$$

In [8], π and f are homogeneous functions (both are convolutions and compute tensors with the same shape), halving ϕ and doubling θ can always lead to decrease in (3), making the lasso loss meaningless. [9] regards \mathbf{g} , where $\mathbf{g} \in \mathbb{R}^{1 \times C'}$, as a replacement of the BN scaling factors, so the following BN layer is modified by eliminating the scaling parameters. In such case, function π need to be carefully designed and trained due to the *internal covariate shift* problem [21] (will be analyzed in the next section).

Different from the above mentioned dynamic execution methods that skips computations for some output cells, the dynamic mechanism in DGC conducts \mathbf{g} on the input channels to sparsify their connections with the output channels, while keeping the shape of the output volume unchanged. Embedding into group executions, DGC actually performs like the multi-head self-attention mechanism proposed in [46]. Thereby, the Eq. 2 is modified as:

$$\hat{f}(\mathbf{x}^i) = f(\mathbf{x}^i \cdot \pi(\mathbf{x}^i, \phi^i), \theta^i). \quad (4)$$

In this case, the scale-sensitive problem can be easily removed by connecting an untouched BN layer (see Section 3.2). Note that in practical computation, θ^i is also partially selected according to the selected channels in \mathbf{x}^i (see Eq. 7).

3.2 Group-wise Dynamic Execution

An illustration of the framework of a DGC layer can be seen in Fig. 2. We split the output channels into multiple groups, each of them is generated by an auxiliary head that equips with an input channel selector to decide which part of input channels should be selected for convolution calculation (see the blue and green areas in Fig. 2). First, the input channel selector in each head adopts a gating strategy to dynamically determine the most important subset of input channels according to their importance scores generated by a saliency generator. Then, the normal convolution is conducted based on the selected subset of input channels generating the output channels in each head. Finally, the output channels from different heads are concatenated and shuffled, which would be connected to a BN layer and non-linear activation layer.

Saliency Generator. The saliency generator assigns each input channel a score representing its importance. Each head has a specific saliency generator, which encourages different heads to use different subpart of the input channels and achieve diversified feature representations. In this paper, we follow the design of the SE block in [16] to design the saliency generator. For the i th head, the saliency vector \mathbf{g}^i is calculated as:

$$\mathbf{g}^i = \pi(\mathbf{x}^i, \phi^i) = \pi(\mathbf{x}, \phi^i) = (W^i(p(\mathbf{x})) + \beta^i)_+, \quad (5)$$

where $\mathbf{g}^i \in \mathbb{R}^{1 \times C}$ represents the saliency vector for the input channels, $(z)_+$ denotes the ReLU activation, p reduces each feature map in \mathbf{x} into a single scalar, such as global average pooling as we used in our experiments, β^i and W^i are trainable parameters representing the biases and a two-step transformation function mapping $\mathbb{R}^{1 \times C} \mapsto \mathbb{R}^{1 \times C/d} \mapsto \mathbb{R}^{1 \times C}$ with d being the squeezing rate. Note that \mathbf{x}^i in eq. 4 is equal to the whole input volume \mathbf{x} here, meaning all the input channels would be considered as candidates in each of the heads.

Gating Strategy. Once the saliency vector is obtained, the next step is to determine which part of the input channels should join the following convolution in the current head. We can either use a head-wise or globally network-wise threshold that decides a certain number of passed gates out of all for each head in a DGC layer. We adopt the head-wise threshold here for simplicity and the global threshold will be discussed in 4.3.

Given a target pruning rate ξ , the head-wise threshold τ^i in i th head meets the equation where $|\mathcal{S}|$ means the length of set \mathcal{S} :

$$\xi = \frac{|\{g^i \mid g^i < \tau^i, g^i \in \mathbf{g}^i\}|}{|\{g^i \mid g^i \in \mathbf{g}^i\}|}. \quad (6)$$

The saliency goes with two streams, *i.e.*, any of channels in the input end with its saliency smaller than the threshold is screened out, while anyone in the remaining channels is amplified with its corresponding g , leading to a group of selected emphasized channels $\mathbf{y}^i \in \mathbb{R}^{(1-\xi)C \times H \times W}$. Assuming the number of heads is \mathcal{H} , the convolution in the i th head is conducted with \mathbf{y}^i and the corresponding filters \mathbf{w}^i ($\mathbf{w}^i \subset \boldsymbol{\theta}^i$, $\boldsymbol{\theta}^i \in \mathbb{R}^{k \times k \times C \times \frac{C'}{\mathcal{H}}}$ and k is the kernel size):

$$\begin{aligned} \hat{\mathbf{x}}^i &= \mathbf{x}[\mathbf{v}, :, :], \quad \hat{\mathbf{g}}^i = \mathbf{g}^i[\mathbf{v}], \quad \mathbf{w}^i = \boldsymbol{\theta}^i[:, :, \mathbf{v}, :], \quad \mathbf{v} = \mathcal{I}_{top}[(1-\xi)C](\mathbf{g}^i) \\ \mathbf{x}'^i &= \hat{f}(\hat{\mathbf{x}}^i) = f(\underbrace{\mathbf{x} \cdot \pi(\mathbf{x}, \phi^i)}_{\text{before pruning}}, \boldsymbol{\theta}^i) = f(\underbrace{\hat{\mathbf{x}}^i \cdot \hat{\mathbf{g}}^i}_{\text{after pruning}}, \mathbf{w}^i) = f(\mathbf{y}^i, \mathbf{w}^i) = \mathbf{y}^i \otimes \mathbf{w}^i, \end{aligned} \quad (7)$$

where $\mathcal{I}_{top}[k](\mathbf{z})$ returns the indices of the k largest elements in \mathbf{z} , the output $\mathbf{x}'^i \in \mathbb{R}^{\frac{C'}{\mathcal{H}} \times H' \times W'}$, \otimes means the regular convolution. At the end of the DGC layer, individual outputs gathered from multiple heads are simply concatenated and shuffled, producing \mathbf{x}' . The gating strategy is unbiased for weights updating since all the channels own the equal chance to be selected.

To induce sparisty, we also add a lasso loss following Eq. 3 as

$$L_{lasso} = \lambda \frac{1}{\mathcal{L}\mathcal{H}} \sum_{l=1}^{\mathcal{L}} \sum_{i=1}^{\mathcal{H}} \|\mathbf{g}^{l,i}\|_1 \quad (8)$$

to the total loss, where \mathcal{L} is the number of DGC layers and λ is a pre-defined parameter.

Computation Cost. Regular convolution using the weight tensor $\boldsymbol{\theta}$ with kernel size k takes $k^2 C' C H W'$ multiply-accumulate operations (MACs). In a DGC layer for each head, the saliency generator and convolution part costs $\frac{2C^2}{d}$ and $k^2(1-\xi)C \frac{C'}{\mathcal{H}} H' W'$ MACs respectively. Therefore, the overall MAC saving of a DGC layer is:

$$(k^2 C C' H' W') / (\mathcal{H}(k^2(1-\xi)C \frac{C'}{\mathcal{H}} H' W' + \frac{2C^2}{d})) = 1 / ((1-\xi) + \frac{2\mathcal{H}C}{dk^2 C' H' W'}) \approx 1 / (1-\xi).$$

As a result, the number of heads \mathcal{H} has negligible influence on the total computation cost.

Invariant to Scaling. As discussed in Section 3.1, [9] implements a convolution layer (including normalization and non-linear activation) by replacing BN scaling factors with the saliency vector:

$$\mathfrak{L}(\mathbf{x}) = (\mathbf{g} \cdot \text{norm}(\text{conv}(\mathbf{x}, \boldsymbol{\theta})) + \boldsymbol{\beta})_+,$$




Fig. 3. The training process of a DGC network. We gradually increase the number of deactivated input channels in each DGC layer with three stages, using a cosine shape learning rate, to make the first stage (first 1/12 of the training epochs, where no channels are pruned) warm up the network, the second stage perform gradual pruning, and the last stage (last 1/4 epochs) fine-tune the sparsified structure. The right figure shows an example of training process on the CIFAR-10 dataset, with pruning rate 0.75 and total 150 epochs. The blue curve shows that even after entering the second stage, the training loss smoothly drops towards 0.

where $\text{norm}(z)$ normalizes each output channel with mean 0 and variance 1. Since $\mathbf{g} = \pi(\mathbf{x}, \phi)$ is dynamically generated, scaling in ϕ first leads to scale change in \mathbf{g} and then in $\mathcal{L}(\mathbf{x})$, thus each sample has a self-dependent distribution, leading to the *internal covariate shift* problem during training on the whole dataset [21]. In contrast, the inference with a DGC layer is defined as (we denote \mathbf{y}^i as \mathbf{y} and \mathbf{w}^i as \mathbf{w} for convenience):

$$\mathcal{L}(\mathbf{x}) = (\text{batchnorm}(\text{conv}(\mathbf{y}, \mathbf{w})))_+.$$

Combining Eq. 7 and above equation, scaling in ϕ leads to scale change in \mathbf{g} while stops at \mathbf{y} . With the following batch normalization, the ill effects of the internal covariate shift can be effectively removed [21].

Training DGC Networks. We train our DGC network from scratch by an end-to-end manner, without the need of model pre-training. During backward propagation, for θ which accompanies a decision process, gradients are calculated only for weights connected to selected channels during the forward pass, and safely set as 0 for others thanks to the unbiased gating strategy. To avoid abrupt changes in training loss while pruning, we gradually deactivate input channels along the training process with a cosine shape learning rate (see Fig. 3).

4 Experiments

4.1 Experimental settings

We evaluate the proposed DGC on three common image classification benchmarks: CIFAR-10, CIFAR-100 [24], and ImageNet (ILSVRC2012) [5]. The two CIFAR datasets both include 50K training images and 10K test images with 10 classes for CIFAR-10 and 100 classes for CIFAR-100. The ImageNet dataset consists of 1.28 million training images and 50K validation images from 1000




Fig. 4. Comparison with SGC and FLGC based on MobileNetV2 on CIFAR-10. For DGC, the width multiplier and the positions of the downsample layers are adjusted to obtain models with matched MACs comparing with SGC and FLGC. For DGC-1 and DGC-2, the downsample layers are set on the first layer of the 3rd and 5th block. While the first layer of the 2nd and 4th block are set as the downsample layers for DGC-3. DGC-1 sets the width multiplier to be 1, while both DGC-2 and DGC-3 set the width multiplier to be 0.75. The results of SGC and FLGC are copied from the FLGC paper [48].

classes. For evaluations, we report detailed accuracy/MACs trade-off against different methods on both CIFAR and ImageNet dataset.

We choose previous popular models including ResNet [11], CondenseNet [19] and MobileNetV2 [44] as the baseline models and embed our proposed DGC into them for evaluations.

ResNet with DGC. ResNet is one of the most powerful models that has been extensively applied in many visual tasks. In this paper, we use the ResNet18 as the baseline model and replace the two 3×3 convolution layers in each residual block with the proposed DGC.

MobileNetV2 with DGC. MobileNetV2 [44] is a powerful efficient network architecture, which is designed for mobile applications. To further improve its efficiency, we replace the last 1×1 point-wise convolution layer in each inverted residual block with our proposed DGC.

CondenseNet with DGC. CondenseNet is also a popular efficient model, which combines the dense connectivity in DenseNet [20] with a learned group convolution (LGC) module. We compare against the LGC by replacing it with the proposed DGC.

All experiments are conducted using Pytorch [39] deep learning library. For both datasets, the standard data augmentation scheme is adopted as in [19]. All models are optimized using stochastic gradient descent (SGD) with Nesterov momentum with a momentum weight of 0.9, and the weight decay is set to be 10^{-4} . Mini batch size is set as 64 and 256 for CIFAR and ImageNet, respectively. The cosine shape learning rate (Fig. 3) starts from 0.1 for ResNet/CondenseNet and 0.05 for MobileNetV2. Models are trained for 300 epochs on CIFAR and 150




Fig. 5. Comparison with SGC and LGC based on CondenseNet on CIFAR-10 (left) and CIFAR-100 (right). The groups (heads) is set as 4 for both LGC and DGC. We change the pruning rate ξ in DGC, *condensation factor* C in LGC and the number of groups G in SGC to control the model computation costs. Noting that for LGC and SGC, when $C = G = 1$, the two models actually share the same structure, corresponding to our DGC version when $\xi = 0$. While a DGC layer works in a self-attention way with \mathcal{H} auxiliary saliency generators, making it better than its vanilla counterparts.

epochs on ImageNet for MobileNetV2, 150 epochs on CIFAR and 120 epochs on ImageNet for CondenseNet, and 120 epochs for ResNet on ImageNet. By default we use 4 groups (heads) in each DGC layer.

4.2 Results on Cifar

We compare DGC with both SGC and previous state-of-the-art learnable group convolutions including CondenseNet [19] and FLGC [48] on CIFAR datasets to demonstrate the effectiveness of its dynamic selecting mechanism. We use MobileNetV2 as backbone and conduct comparison with SGC and FLGC. For comparison with CondenseNet, we employ the proposed modified DenseNet structure in CondenseNet with 50 layers as the baseline model.

MobileNetV2. We compare the proposed DGC with SGC and FLGC on CIFAR-10. Like the FLGC, we replace the last 1×1 convolution layer in each inverted residual block with the proposed DGC. For SGC and FLGC, models with different MACs are obtained by adjusting the number of groups in group convolution layers. For our proposed DGC, we adjust the width multiplier and the positions of the downsample convolution layers (*i.e.*, with stride equals to 2) to obtain models with matched MACs. The pruning rate ξ is set to be 0.65 for each head and 4 heads are employed in a DGC layer. The results are shown in Fig. 4. It can be seen that both DGC and FLGC outperforms SGC, while DGC can achieve lowest top-1 errors with less computation costs. Note that the group structures are still fixed in FLGC after training, which ignores the properties of single inputs and results in poor connections for some of them. By introducing dynamic feature selectors, our proposed DGC can adaptively select

Table 1. Comparison of Top-1 and Top-5 classification error with state-of-the-art filter pruning and dynamic channel selection methods using ResNet-18 as the baseline model on ImageNet.

Model	Group	Dynamic	Top-1	Top-5	MAC Saving
SFP [12]			32.90	12.22	(1.72×)
NS [35]			32.79	12.61	(1.39×)
DCP [57]			32.65	12.40	(1.89×)
FPGM [13]			31.59	11.52	(1.53×)
LCCN [8]		✓	33.67	13.06	(1.53×)
FBS [9]		✓	31.83	11.78	(1.98×)
CGNet [17]	✓	✓	31.70	-	(2.03×)
DGC	✓	✓	31.22	11.38	(2.04×)
DGC-G	✓	✓	31.37	11.56	(2.08×)

most related input channels for each group conditioned on the individual input images. The results indeed demonstrate the superiority of DGC compared with FLGC. Specifically, DGC achieves lower top-1 errors and computation cost (for example, 5.17% *vs.* 5.89% and 155M *vs.* 158M).


CondenseNet. We compare the proposed DGC with SGC and CondenseNet on both CIFAR-10 and CIFAR-100. CondenseNet adopts a LGC strategy to sparsify the compute-intensive 1x1 convolution layers and uses the group-lasso regularizer to gradually prune less important connections during training. The final model after training is also fixed like FLGC. We replace all the LGC layers in CondenseNet with the SGC or DGC structures under similar computation costs for comparison. The results are shown in Fig. 5. It also shows both DGC and LGC perform better than SGC due to the learnable soft channel assignments for groups. Meanwhile, DGC can perform better than LGC with less or similar MACs (for example, 5.83% *vs.* 6.10% and 51M *vs.* 66M on CIFAR-10, 26.18% *vs.* 26.58% and 51M *vs.* 66M on CIFAR-100), demonstrating that the dynamic computation scheme in DGC is superior to the static counterpart in LGC.

4.3 Results on ImageNet

To further demonstrate the effectiveness of the proposed DGC, we compare it with state-of-the-art filter pruning and channel-level dynamic execution methods using ResNet18 as the baseline model on ImageNet. We also embed the DGC into the CondenseNet [19] and MobileNetV2 [44] (*i.e.*, CondenseNet-DGC and MobileNetV2-DGC) to compare with state-of-the-art efficient CNNs. For CondenseNet-DGC/SGC, the same structure as CondenseNet in [19] is used with LGC replaced by DGC/SGC. For MobileNetV2-DGC, the last 1×1 convolution layers in all inverted residual blocks in the original MobileNetV2 are replaced with DGC layers. We set $\mathcal{H} = 4$, $\xi = 0.75$ and $\lambda = 10^{-5}$ for all models. The results are reported in Table 1 and Table 2. Here, for further exploration, we additionally adopt a global threshold (see appendix for detail) to give more flexibility to the DGC network. The corresponding models are noted with a suffix “G”. Different from the head-wise threshold, the global threshold is learnt

Table 2. Comparison of Top-1 and Top-5 classification error with state-of-the-art efficient CNNs on ImageNet.

Model	Top-1	Top-5	MACs
MobileNetV1 [15]	29.4	10.5	569M
ShuffleNet [55]	29.1	10.2	524M
NASNet-A (N=4) [58]	26.0	8.4	564M
NASNet-B (N=4) [58]	27.2	8.7	488M
NASNet-C (N=3) [58]	27.5	9.0	558M
IGCV3-D [45]	27.8	-	318M
MobileNetV2 [44]	28.0	9.4	300M
CondenseNet [19]	26.2	8.3	529M
CondenseNet-SGC	29.0	9.9	529M
CondenseNet-FLGC [48]	25.3	7.9	529M
MobileNetV2-DGC	29.3	10.2	245M
CondenseNet-DGC	25.4	7.8	549M
CondenseNet-DGC-G	25.2	7.8	543M

**Fig. 6.** Ablation study on head numbers and pruning rates

during training and used in testing (like parameters in BN layers), therefore the actual pruning rate is slightly different to the target ξ .

Table 1 shows that our method outperforms all of the compared static filter pruning methods while achieving higher speed-up. For example, when compared with previous best-performed FPGM [13] that achieves $1.53\times$ speed-up, our method decreases the top-1 error by 0.37% with $2.04\times$ speed-up. It can also be seen that our method is superior to those dynamic execution methods including LCCN [8], FBS [9], and CGNet [17]. Compared with FBS, our DGC can better achieve diversified feature representations with multiple heads that work in a self-attention way. As for LCCN and CGNet, both of them work in a pixel-level dynamic execution way, which results in irregular computations and needs specific algorithms or devices for computation acceleration. However, our DGC directly removes some part of input channels for each group, which can be easily implemented like SGC.

It can be found from Table 2 that the proposed DGC performs well when embedded into the state-of-the-art efficient CNNs. Specifically, DGC can further speed up the MobileNetV2 by $1.22\times$ (245M *vs.* 300M) with only 1.3% degradation of top-1 error. In addition, DGC outperforms both SGC




Fig. 7. Visualization of network sparse patterns for different input images. Each row represents a particular image. We compare the saliency vectors in shallow (*the 5-8th DGC layers*) and deep (*the 33-36th DGC layers*) DGC layers of the CondenseNet-DGC structure used in 4.3 and also output their corresponding pruning decisions as white-black pixels. In each saliency map (with 64 columns and 16 rows), columns mean different channels, each row indicates a certain head for a layer and every 4 rows represent a DGC layer since $\mathcal{H} = 4$.

and CondenseNet’s LGC, and obtains comparable accuracy when compared with FLGC. Finally, our MobileNetV2-DGC even achieves comparable accuracy compared with MobileNetV1 and ShuffleNet while reducing the computation cost more than doubled.

4.4 Ablation and Visualization

Number of Heads. We use a 50-layer CondenseNet-DGC network structure on the CIFAR-10 dataset as the basis to observe the effects of number of heads ($\mathcal{H} = 1, 2, 4, 8$). The network is trained for 150 epochs and the pruning rate ξ is set to 0.75. Blue bars in Fig. 6 shows the Top-1 error under varying head numbers. It can be seen that the performance is firstly improved with the increase of head number, reaching the peak at $\mathcal{H} = 4$, and then slightly drops when $\mathcal{H} = 8$. We conjecture that under similar computation costs, using more heads helps capture diversified feature representations (also see Fig. 8), while a big head number may also bring extra difficulty in training effectiveness.

Pruning Rate. We follow the above structure but fix $\mathcal{H} = 4$ and adjust the pruning rate ($\xi = 7/8, 5/6, 3/4, 1/2, 1/3$). The right part in Fig. 6 shows the performance changes. Generally, with a higher pruning rate, more channels are deactivated during inference thus with less computation cost, but at the same time leading to an increasing error rate.

Dynamicness and Adaptiveness of DGC networks. Fig. 7 visualizes the saliency vectors and corresponding sparse patterns of the CondenseNet-DGC structure used in 4.3 for different input images. We can see that shallower layers share the similar sparse patterns, since they tend to catch the basic and less abstract image features, while deeper layers diverge different images into different sparse patterns, matching our observation in Section 1. On the other hand, similar images tend to produce similar patterns, and vice versa, which indicates the *adaptiveness* property of DGC networks.

In addition, We also visualize the pruning probability of channels in one of the deep layers in Fig. 8. We track the first two heads and show the *jumping* channels




Fig. 8. Channel pruning probabilities. We track the last 900 input channels in the 33rd DGC layer of the CondenseNet-DGC structure used in 4.3 on the validation set of ImageNet dataset for the first two (out of four) heads (blue and red colors). For each channel each head, the probability is calculated as the percentage of images which deactivate the channel in this head during their inferences.

whose pruning probability is neither 0 nor 1 (noting that those *freezing* channels with probability 0 or 1 can still be jumping in the other two heads). Firstly, We can see that each head owns a particular pruning pattern but complementary with the other, indicating features can be captured from multiple perspectives with different heads. Secondly, the existence of *jumping* channels reveals the *dynamicness* of the DGC network, which will adaptively ignite its channels depending on certain input samples (more results can be seen in appendix).

Computation time. Following [19,48], the inference speed of DGC can be made close to SGC by embedding a dynamic index layer right before each convolution that reorders the indices of input channels and filters according to the group information. Based on this, we test the speed of DGC on Resnet18 with an Intel i7-8700 CPU by calculating the average inference time spent on the convolutional layers, using the original model (11.5ms), SGC with $G = 4$ (5.1ms), and DGC with $\xi = 3/4$ and $\mathcal{H} = 4$ (7.8ms = 5.1ms + 0.8ms + 1.9ms). Specifically, in DGC, the saliency generator takes 0.8ms and the dynamic index layer takes 1.9ms, the rest is the same as of SGC. However, the speed can be further improved with a more careful optimization.

5 Conclusion

In this paper, we propose a novel group convolution method named DGC, which improves the existing group convolutions that have fixed connections during inference by introducing the dynamic channel selection scheme. The proposed DGC has the following three-fold advantages over existing group convolutions and previous state-of-the-art channel/pixel-level dynamic execution methods: 1) It conducts sparse group convolution operations while keeping the capabilities of the original CNNs; 2) It dynamically executes sample-dependent convolutions with multiple complementary heads using a self-attention based decision strategy; 3) It can be embedded into various exiting CNNs and the resulted models can be stably and easily optimized in an end-to-end manner without the need of pre-training.

References

1. Banner, R., Nahshan, Y., Soudry, D.: Post training 4-bit quantization of convolutional networks for rapid-deployment. In: *Advances in Neural Information Processing Systems*. pp. 7948–7956 (2019) [4](#)
2. Bolukbasi, T., Wang, J., Dekel, O., Saligrama, V.: Adaptive neural networks for efficient inference. In: *Proceedings of the 34th International Conference on Machine Learning-Volume 70*. pp. 527–536. *JMLR. org* (2017) [4](#)
3. Cao, S., Ma, L., Xiao, W., Zhang, C., Liu, Y., Zhang, L., Nie, L., Yang, Z.: SeerNet: Predicting convolutional neural network feature-map sparsity through low-bit quantization. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 11216–11225 (2019) [4](#)
4. Chollet, F.: Xception: Deep learning with depthwise separable convolutions. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 1251–1258 (2017) [2](#)
5. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: *2009 IEEE conference on computer vision and pattern recognition*. pp. 248–255. *Ieee* (2009) [1](#), [8](#)
6. Denton, E.L., Zaremba, W., Bruna, J., LeCun, Y., Fergus, R.: Exploiting linear structure within convolutional networks for efficient evaluation. In: *Advances in neural information processing systems*. pp. 1269–1277 (2014) [4](#)
7. Ding, Y., Liu, J., Xiong, J., Shi, Y.: On the universal approximability and complexity bounds of quantized relu neural networks. *arXiv preprint arXiv:1802.03646* (2018) [4](#)
8. Dong, X., Huang, J., Yang, Y., Yan, S.: More is less: A more complicated network with less inference complexity. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 5840–5848 (2017) [4](#), [5](#), [6](#), [11](#), [12](#)
9. Gao, X., Zhao, Y., Dudziak, L., Mullins, R., Xu, C.z.: Dynamic channel pruning: Feature boosting and suppression. *arXiv preprint arXiv:1810.05331* (2018) [3](#), [4](#), [5](#), [6](#), [7](#), [11](#), [12](#)
10. Han, S., Pool, J., Tran, J., Dally, W.: Learning both weights and connections for efficient neural network. In: *Advances in neural information processing systems*. pp. 1135–1143 (2015) [4](#)
11. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 770–778 (2016) [1](#), [3](#), [9](#)
12. He, Y., Kang, G., Dong, X., Fu, Y., Yang, Y.: Soft filter pruning for accelerating deep convolutional neural networks. *arXiv preprint arXiv:1808.06866* (2018) [11](#)
13. He, Y., Liu, P., Wang, Z., Hu, Z., Yang, Y.: Filter pruning via geometric median for deep convolutional neural networks acceleration. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 4340–4349 (2019) [4](#), [11](#), [12](#)
14. Heo, B., Kim, J., Yun, S., Park, H., Kwak, N., Choi, J.Y.: A comprehensive overhaul of feature distillation. In: *Proceedings of the IEEE International Conference on Computer Vision*. pp. 1921–1930 (2019) [4](#)
15. Howard, A.G., Zhu, M., Chen, B., Kalenichenko, D., Wang, W., Weyand, T., Andreetto, M., Adam, H.: Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861* (2017) [3](#), [12](#)
16. Hu, J., Shen, L., Sun, G.: Squeeze-and-excitation networks. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 7132–7141 (2018) [6](#)

17. Hua, W., Zhou, Y., De Sa, C.M., Zhang, Z., Suh, G.E.: Channel gating neural networks. In: *Advances in Neural Information Processing Systems*. pp. 1884–1894 (2019) [3](#), [4](#), [5](#), [11](#), [12](#)
18. Huang, G., Chen, D., Li, T., Wu, F., van der Maaten, L., Weinberger, K.Q.: Multi-scale dense networks for resource efficient image classification. *arXiv preprint arXiv:1703.09844* (2017) [4](#)
19. Huang, G., Liu, S., Van der Maaten, L., Weinberger, K.Q.: Condensenet: An efficient densenet using learned group convolutions. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 2752–2761 (2018) [2](#), [3](#), [5](#), [9](#), [10](#), [11](#), [12](#), [14](#)
20. Huang, G., Liu, Z., Van Der Maaten, L., Weinberger, K.Q.: Densely connected convolutional networks. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 4700–4708 (2017) [1](#), [2](#), [3](#), [9](#)
21. Ioffe, S., Szegedy, C.: Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167* (2015) [5](#), [6](#), [8](#)
22. Jaderberg, M., Vedaldi, A., Zisserman, A.: Speeding up convolutional neural networks with low rank expansions. *arXiv preprint arXiv:1405.3866* (2014) [4](#)
23. Jin, X., Peng, B., Wu, Y., Liu, Y., Liu, J., Liang, D., Yan, J., Hu, X.: Knowledge distillation via route constrained optimization. In: *Proceedings of the IEEE International Conference on Computer Vision*. pp. 1345–1354 (2019) [4](#)
24. Krizhevsky, A., Hinton, G., et al.: Learning multiple layers of features from tiny images (2009) [8](#)
25. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. In: *Advances in neural information processing systems*. pp. 1097–1105 (2012) [2](#), [3](#)
26. Lee, N., Ajanthan, T., Torr, P.H.: Snip: Single-shot network pruning based on connection sensitivity. *arXiv preprint arXiv:1810.02340* (2018) [4](#)
27. Li, Y., Lin, S., Zhang, B., Liu, J., Doermann, D., Wu, Y., Huang, F., Ji, R.: Exploiting kernel sparsity and entropy for interpretable cnn compression. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 2800–2809 (2019) [4](#)
28. Liang, X.: Learning personalized modular network guided by structured knowledge. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 8944–8952 (2019) [4](#)
29. Lin, J., Rao, Y., Lu, J., Zhou, J.: Runtime neural pruning. In: *Advances in Neural Information Processing Systems*. pp. 2181–2191 (2017) [3](#), [4](#), [5](#)
30. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: *European conference on computer vision*. pp. 740–755. Springer (2014) [1](#)
31. Liu, C., Ding, W., Xia, X., Zhang, B., Gu, J., Liu, J., Ji, R., Doermann, D.: Circulant binary convolutional networks: Enhancing the performance of 1-bit dcnn with circulant back propagation. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 2691–2699 (2019) [4](#)
32. Liu, J., Wen, D., Gao, H., Tao, W., Chen, T.W., Osa, K., Kato, M.: Knowledge representing: Efficient, sparse representation of prior knowledge for knowledge distillation. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*. pp. 0–0 (2019) [4](#)
33. Liu, L., Deng, J.: Dynamic deep neural networks: Optimizing accuracy-efficiency trade-offs by selective execution. In: *Thirty-Second AAAI Conference on Artificial Intelligence* (2018) [4](#)

34. Liu, Y., Dong, W., Zhang, L., Gong, D., Shi, Q.: Variational bayesian dropout with a hierarchical prior. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 7124–7133 (2019) [4](#)
35. Liu, Z., Li, J., Shen, Z., Huang, G., Yan, S., Zhang, C.: Learning efficient convolutional networks through network slimming. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 2736–2744 (2017) [11](#)
36. Long, J., Shelhamer, E., Darrell, T.: Fully convolutional networks for semantic segmentation. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 3431–3440 (2015) [1](#)
37. Ma, N., Zhang, X., Zheng, H.T., Sun, J.: Shufflenet v2: Practical guidelines for efficient cnn architecture design. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 116–131 (2018) [2](#), [3](#)
38. Minnehan, B., Savakis, A.: Cascaded projection: End-to-end network compression and acceleration. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 10715–10724 (2019) [4](#)
39. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al.: Pytorch: An imperative style, high-performance deep learning library. In: Advances in Neural Information Processing Systems. pp. 8024–8035 (2019) [9](#)
40. Peng, H., Wu, J., Chen, S., Huang, J.: Collaborative channel pruning for deep networks. In: International Conference on Machine Learning. pp. 5113–5122 (2019) [4](#)
41. Phuong, M., Lampert, C.: Towards understanding knowledge distillation. In: International Conference on Machine Learning. pp. 5142–5151 (2019) [4](#)
42. Phuong, M., Lampert, C.H.: Distillation-based training for multi-exit architectures. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 1355–1364 (2019) [4](#)
43. Sakr, C., Shanbhag, N.: Per-tensor fixed-point quantization of the back-propagation algorithm. arXiv preprint arXiv:1812.11732 (2018) [4](#)
44. Sandler, M., Howard, A., Zhu, M., Zhmoginov, A., Chen, L.C.: Mobilenetv2: Inverted residuals and linear bottlenecks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 4510–4520 (2018) [3](#), [9](#), [11](#), [12](#)
45. Sun, K., Li, M., Liu, D., Wang, J.: Igc3: Interleaved low-rank group convolutions for efficient deep neural networks. BMVC (2018) [2](#), [3](#), [12](#)
46. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. In: Advances in neural information processing systems. pp. 5998–6008 (2017) [6](#)
47. Veit, A., Belongie, S.: Convolutional networks with adaptive inference graphs. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 3–18 (2018) [4](#)
48. Wang, X., Kan, M., Shan, S., Chen, X.: Fully learnable group convolution for acceleration of deep neural networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 9049–9058 (2019) [2](#), [3](#), [5](#), [9](#), [10](#), [12](#), [14](#)
49. Wang, X., Yu, F., Dou, Z.Y., Darrell, T., Gonzalez, J.E.: Skipnet: Learning dynamic routing in convolutional networks. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 409–424 (2018) [4](#)
50. Wu, Z., Nagarajan, T., Kumar, A., Rennie, S., Davis, L.S., Grauman, K., Feris, R.: Blockdrop: Dynamic inference paths in residual networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 8817–8826 (2018) [4](#)

51. Xiao, X., Wang, Z., Rajasekaran, S.: Autoprune: Automatic network pruning by regularizing auxiliary parameters. In: *Advances in Neural Information Processing Systems*. pp. 13681–13691 (2019) [4](#)
52. Xie, S., Girshick, R., Dollár, P., Tu, Z., He, K.: Aggregated residual transformations for deep neural networks. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 1492–1500 (2017) [2](#), [3](#)
53. Yoo, J., Cho, M., Kim, T., Kang, U.: Knowledge extraction with no observable data. In: *Advances in Neural Information Processing Systems*. pp. 2701–2710 (2019) [4](#)
54. Zhang, T., Qi, G.J., Xiao, B., Wang, J.: Interleaved group convolutions. In: *Proceedings of the IEEE international conference on computer vision*. pp. 4373–4382 (2017) [3](#)
55. Zhang, X., Zhou, X., Lin, M., Sun, J.: Shufflenet: An extremely efficient convolutional neural network for mobile devices. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 6848–6856 (2018) [2](#), [3](#), [12](#)
56. Zhang, X., Zou, J., Ming, X., He, K., Sun, J.: Efficient and accurate approximations of nonlinear convolutional networks. In: *Proceedings of the IEEE Conference on Computer Vision and pattern Recognition*. pp. 1984–1992 (2015) [4](#)
57. Zhuang, Z., Tan, M., Zhuang, B., Liu, J., Guo, Y., Wu, Q., Huang, J., Zhu, J.: Discrimination-aware channel pruning for deep neural networks. In: *Advances in Neural Information Processing Systems*. pp. 875–886 (2018) [11](#)
58. Zoph, B., Vasudevan, V., Shlens, J., Le, Q.V.: Learning transferable architectures for scalable image recognition. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 8697–8710 (2018) [12](#)