

CompConv: A Compact Convolution Module for Efficient Feature Learning

Chen Zhang¹ Yinghao Xu² Yujun Shen²

¹Zhejiang University ²The Chinese University of Hong Kong

daisy_chen@zju.edu.cn {xy119, sy116}@ie.cuhk.edu.hk

Abstract

Convolutional Neural Networks (CNNs) have achieved remarkable success in various computer vision tasks but rely on tremendous computational cost. To solve this problem, existing approaches either compress well-trained large-scale models or learn lightweight models with carefully designed network structures. In this work, we make a close study of the convolution operator, which is the basic unit used in CNNs, to reduce its computing load. In particular, we propose a compact convolution module, called *CompConv*, to facilitate efficient feature learning. With the divide-and-conquer strategy, *CompConv* is able to save a great many computations as well as parameters to produce a certain dimensional feature map. Furthermore, *CompConv* discreetly integrates the input features into the outputs to efficiently inherit the input information. More importantly, the novel *CompConv* is a plug-and-play module that can be directly applied to modern CNN structures to replace the vanilla convolution layers without further effort. Extensive experimental results suggest that *CompConv* can adequately compress the benchmark CNN structures yet barely sacrifice the performance, surpassing other competitors.

1. Introduction

In recent years, Convolutional Neural Networks (CNNs) have significantly advanced many tasks in the computer vision field [11, 21, 36] due to their great power in learning representative features. However, such success relies on massive computation and storage resources, hindering these large-scale models from being deployed on resource-limited devices in practice.

To alleviate the pressure on the demand for computing power and memory, one straight-forward solution is to compress well-trained big models. For this purpose, existing approaches propose either to employ a tiny model as the student to distill the knowledge from the teacher model [13, 8], or to prune the unimportant connections from the learned model to make it thinner [22, 12, 10].

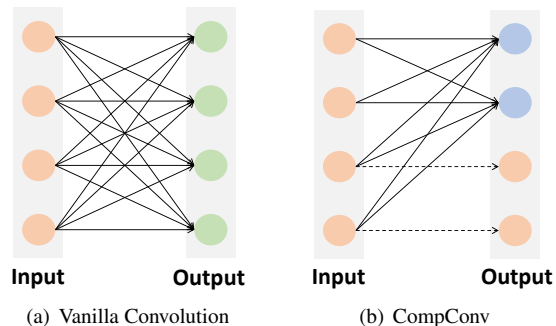


Figure 1. Concept diagram of the proposed compact convolution module, *CompConv*. Compared to the vanilla convolution in (a), which fully connects the input channels with the output channels, *CompConv* adopts the divide-and-conquer strategy to develop the output recursively. (b) gives a brief illustration of the core unit used in *CompConv*, which only uses the convolution operation to contribute half number of channels (*i.e.*, the solid line) and subtly borrows the other half from the inputs (*i.e.*, the dashed line).

Besides training large-scale models and then performing compression, an alternative solution is to directly learn a lightweight model. Nevertheless, a side effect of reducing the model size is the sacrifice of the learning capacity. To tackle this obstacle, it often requires careful design of the network architecture, such as MobileNet [15, 35] and SqueezeNet [18].

In this paper, different from prior work, we take a close look into the convolution operator, which is the basic unit used in all CNNs, and propose to save the computing load by compressing the convolution module. Recall that the convolution module learns the transformation from one feature space to another. More concretely, given an input feature map, with C_{in} channels, the convolutional kernel projects it to a C_{out} -channels feature map as the output. For each pixel in the spatial field, the conventional convolution requires $C_{in} \times C_{out}$ connections, as shown in Fig. 1a. As pointed out by prior work [9, 22, 10], however, the learning of CNNs presents a definite over-parameterization and redundancy. From this perspective, we propose *CompConv* to implement the convolution more compactly.

Specifically, instead of directly producing the final feature map, CompConv employs a core unit that equally splits the output into two parts along the channel axis. One of them is transformed from the input feature via convolutional projection while the other identically borrows a subset of input channels. In this way, we are able to pass down the learned information across the convolution module to the most extent yet with minor effort. Fig. 1b illustrates the diagram of such process. More importantly, our core unit can be performed in a recursive computing manner, resulting in a divide-and-conquer strategy. As a result, CompConv is able to save a great many computations as well as parameters for more efficient feature extraction. We summarize our contributions as follows:

- We propose a compact convolution module, termed as *CompConv*, which utilizes the divide-and-conquer strategy together with the carefully designed identical mapping to considerably reduce the computational cost of CNNs.
- We make exhaustive analysis of the proposed CompConv by studying how the learning capacity is affected by the depth of the recursive computation. We further propose a practical scheme to convincingly control the compression rate.
- We apply CompConv to various benchmark CNNs as a handy replacement of the conventional convolution layer. It turns out that CompConv can substantially save the computing load yet barely sacrifice the model performance on classification and detection tasks, outperforming existing approaches.

2. Related Work

Model Compression. Modern deep neural networks have millions of weights, rendering them both memory-intensive and computationally expensive. A straightforward way to reduce the computational cost is to remove the unimportant neuron connections with rule-based or learning-based methods, named network pruning [10, 22, 12, 31, 30, 6, 31, 24, 40, 17, 23, 26, 25]. Recent work also shows that precision computation is not necessary for the training and inference of deep models [7]. As a result, quantizing weights and activations [19, 44] are also widely explored to improve network efficiency. Specifically, binary networks [4, 33], which only employ 1-bit neuron to represent the model weights and activations, replace all MAC operation with boolean operation to save the computation. Besides, knowledge distillation [13, 8, 41, 16, 39, 39] offers an alternative way to generate small student networks with the guidance of a well-trained large teacher network. In this way, the students show comparable performance as the teacher but raise a more efficient inference process.

Lightweight Model Design. Another way of reducing the computational cost of CNNs is to directly learn a lightweight model, but there is a clear trade-off between the model size and the learning capacity. As a result, existing approaches propose to increase the model efficiency and maintain its performance at the same time with carefully designed network structure [18, 15, 35, 42, 32, 14]. SqueezeNet [18] adopts a large amount of 1×1 convolutions to reduce the number of parameters. MobileNet V1 [15] and V2 [35] employ depth-wise separable convolutions and inverted linear residual bottleneck to improve computation efficiency. ShuffleNet V1 [42] and V2 [32] propose the channel-shuffle operation to enhance the information flow between different channel groups and provide a hardware-friendly implementation to enable practical applications. GhostNet [9] considers the feature redundancy between feature maps and proposes to learn ghost features with cheap operations. On the other hand, Network Architecture Search (NAS) [28, 46, 45, 43] aims at finding the most efficient network structure automatically. MobileNet V3 [14] utilizes the Auto-ML technology [1, 2] to achieve better performance with fewer floating point operations (FLOPs).

3. Compact Convolution Module

In this section, we introduce the proposed CompConv module, which efficiently learns the output feature from the input using the divide-and-conquer strategy. Sec. 3.1 introduces the motivation. Sec. 3.2 describes the core unit used in CompConv. Sec. 3.3 presents the complete CompConv module, which executes the core unit in a recursive fashion. Sec. 3.4 provides an adaptive strategy of using CompConv in practice. Sec. 3.5 analyzes the computing complexity of CompConv.

3.1. Motivation

Convolution can be treated as an operation that maps features from one space to another. To some extent, this process is similar to Discrete Fourier Transform (DFT), which maps a signal sequence from the time domain to the frequency domain. Fast Fourier Transform (FFT) is widely used to speed up the computation of DFT. Motivated by FFT, our CompConv is proposed to compress the vanilla convolution module through the divide-and-conquer strategy.¹

Let us review the formulation of FFT. When applying DFT to a N -points signal sequence $x(t)$ from the time domain, FFT proposes to split it into two $\frac{N}{2}$ -points subsequences, denoted as $x^{(e)}(t)$ and $x^{(o)}(t)$, and perform DFT

¹Note that the improvement from DFT to FFT does not change the computational results, *i.e.*, they are identical. Differently, the change from the vanilla convolution to CompConv is not a lossless compression. We simply design CompConv by drawing lessons from FFT for more efficient feature learning.

on each of them. Here, e and o stand for “even” and “odd” respectively. Accordingly, the final result $X(k)$ from the frequency domain can be obtained from the intermediate transformation results $X^{(e)}(k)$ and $X^{(o)}(k)$ with

$$X(k) = X^{(e)}(k) + W_N^k X^{(o)}(k), \quad (1)$$

where $W_N^k = \exp(-j\frac{2\pi}{N}k)$ is a multiplier. Based on this, the factorized results $X^{(e)}(k)$ and $X^{(o)}(k)$ can be further divided into smaller groups, resulting in a recursive computing manner.

3.2. Core Unit of CompConv

Inspired by FFT, we introduce the divide-and-conquer strategy into the convolution module used in CNNs to improve its computing efficiency. By analogy, we treat the intermediate feature maps produced by CNNs as the sequence from the *channel* axis. More concretely, to develop a feature map \mathbf{X} with C channels, we can alternatively develop two feature maps \mathbf{X}_A and \mathbf{X}_B , each of which is with $\frac{C}{2}$ channels, and compose them together with

$$\mathbf{X} = \mathbf{X}_A \oplus \mathbf{W}\mathbf{X}_B, \quad (2)$$

where \oplus denotes the concatenating operation along the channel axis and \mathbf{W} is a learnable parameter used to transform feature maps.

Eq. (2) embodies the key idea of CompConv. In particular, the core unit of CompConv is implemented with two parts, as shown in Fig. 2. One part (*i.e.*, \mathbf{X}_A) is identically mapped from a subset of input channels, which is able to inherit information from the input with minor effort. The other part (*i.e.*, \mathbf{X}_B) is transformed from the input feature with a native convolution module.

3.3. Recursive Computation

With the formulation in Eq. (2), CompConv can be computed in a recursive manner by further splitting \mathbf{X}_B into two halves as

$$\mathbf{X}_{B_i} = \mathbf{X}_{A_{i+1}} \oplus \mathbf{W}_{i+1}\mathbf{X}_{B_{i+1}} \quad i = 0, \dots, d-1, \quad (3)$$

where d denotes the recursion depth.

Tailing Channels. We treat the first separation step (*i.e.*, $\{\mathbf{X}_{A_0}, \mathbf{X}_{B_0}\}$) differently from other steps, as shown in Fig. 2. Concretely, \mathbf{X}_{A_0} is not directly borrowed from the input but transformed from \mathbf{X}_{B_0} instead. There are mainly two reasons in doing so. On one hand, \mathbf{X}_{A_0} is with the most channels among all identical parts $\{\mathbf{X}_{A_i}\}_{i=0}^{d-1}$. If we directly duplicate some input channels as \mathbf{X}_{A_0} , there will be too much redundancy between the input feature map and the output one, severely limiting the learning capacity of this module. On the other hand, besides transformed from \mathbf{X}_{B_0} , there are some alternative ways to obtain \mathbf{X}_{A_0} , such

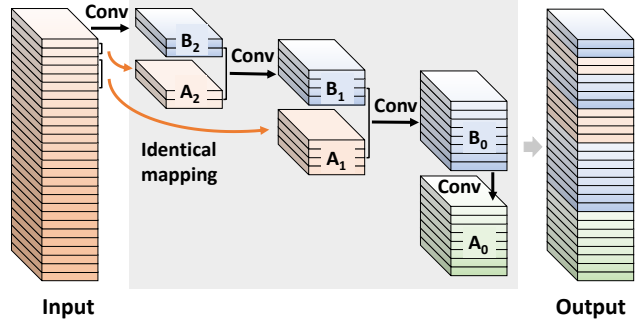


Figure 2. Framework of the proposed CompConv with the recursion depth $d = 3$. With the divide-and-conquer strategy, CompConv computes the output from the input recursively. In each step, to get a feature map with a certain number of channels, the core unit of CompConv proposes to generate half of the channels (blue blocks) with the vanilla convolution operation and the other half (tangerine blocks) with the identical mapping from the input. Meanwhile, to increase the learning capability yet maintain low computing load, the identical mapping is replaced with an additional convolution module in the first separation step to produce tailing channels (green blocks).

as mapping from the entire input feature map or building another recursion like the computation of \mathbf{X}_{B_0} . Among all these approaches, developing \mathbf{X}_{A_0} from \mathbf{X}_{B_0} is with the cheapest computing cost. Meanwhile, the deduction of \mathbf{X}_{B_0} has already assembled enough information from the input feature, hence the learning capacity can also be maintained. We use group convolution with group numbers equal to channel number for this transformation.

Integrating Recursive Results. To better utilize the computations in the recursive process, the final output is developed by not only grouping the two largest sub-features (*i.e.*, \mathbf{X}_{A_0} and \mathbf{X}_{B_0}) but also integrating all the intermediate results, as shown in Fig. 2. In this way, we can make sufficient use of all the computing operations to produce the final output. In addition, a shuffle block is added after the concatenation of these feature maps.

3.4. Adaptive Separation Strategy

As described in Sec. 3.3, CompConv adopts the divide-and-conquer strategy for efficient feature learning. Thus how to recursively split the channels is critical to its computing efficiency and learning capacity. Here, we use C_{in} and C_{out} to denote the number of input channels and output channels respectively. C_{prim} refers to the channel number of the minimum computing unit, *e.g.*, \mathbf{X}_{B_2} when $d = 3$ in Fig. 2. Considering the exponential growth of the channel number during the recursive computation, we can

expect

$$C_{out} = \sum_{i=1}^d 2^i C_{prim}. \quad (4)$$

We can easily solve Eq. (4) with

$$C_{prim} = \lceil \frac{C_{out}}{2 \times (2^d - 1)} \rceil, \quad (5)$$

where $\lceil \cdot \rceil$ denotes the ceiling function that grounds C_{prim} as an integer. If the sum of the channels from all units is larger than C_{out} , we simply drop some last channels in \mathbf{X}_{A_0} to make sure the output feature is with the proper dimension.

Choice of Recursion Depth. We can tell from Eq. (5) that C_{prim} is highly dependent of the recursion depth d , which is a hyper-parameter in our CompConv module. Larger d correspond to higher compression rate, where $d = 0$ means no compression. Considering the different structures and different model scales of modern CNNs, we propose an adaptive strategy on choosing the depth as

$$d = \max(\log_2(\max(1, \frac{C_{in}}{C_0})) + 1, 3). \quad (6)$$

Here, C_0 is a model-specific design choice, valuing from $\{32, 64, 128, 256, 512, \dots\}$, which can be determined by the target compression rate and the model size. Intuitively, larger C_0 tends to induce smaller d , further resulting in a lighter compression. From this perspective, C_0 can be used to control the trade-off between computing efficiency and learning capacity. It's worth noting that the recursion depth d is related to the number of input channels, C_{in} , in Eq. (6), which means our strategy will dynamically adjust the computation depth at different layers. Meanwhile, to guarantee the minimum unit to have adequate learning capacity, we assign it with enough channels. In other words, C_{prim} could not be too small. From Eq. (5), we see that when $d = 3$, C_{prim} only accounts for $\sim 8\%$ of output channels. we therefore bound the depth d with a maximum value 3.

Recommended Configuration. For most popular CNN networks like VGG [36] and ResNet [11], we recommend to set $C_0 = 128$. We denote this setting as CompConv128. In the following sections, we use CompConv128 by default if no special instructions.

3.5. Complexity Analysis

In this part, we briefly discuss how CompConv can help to save the computing load. We conduct the complexity analysis and make a comparison between vanilla convolution and the proposed CompConv. Here, we focus on the number of operations executed in each module.

Assuming both the input and output feature maps are with resolution $H \times W$, the computing complexities for vanilla convolution and CompConv are

$$\mathcal{O}_{Conv} = H \times W \times k^2 \times C_{in} \times C_{out}, \quad (7)$$

$$\mathcal{O}_{CompConv} = H \times W \times k^2 \times (C_{in} \times C_{prim} + \sum_{i=1}^{d-1} (2^i C_{prim})^2 + 2^{d-1} C_{prim}), \quad (8)$$

where k indicates the size of the convolutional kernel. Under the setting $C_{in} = C_{out}$ and $d = 3$, CompConv only requires $\sim 20\%$ computational resources compared to the conventional convolution to develop the output feature with the same number of channels.

4. Experiments

In this section, we evaluate the effectiveness and efficiency of the proposed CompConv on various visual benchmarks, including CIFAR-10 [20], CIFAR-100 [20] and ImageNet [5]. Top-1 and Top-5 accuracy are reported as the evaluation metric for comparison. Ablation studies on various CompConv designs are also included. Besides, detection results on COCO [27] are reported to show the efficiency and generality of CompConv. The standard mean average precision (mAP) is used to measure accuracy.

This section is organized as follows: Sec. 4.1 introduces the experiment setting, including datasets and the details of training and inference. Sec. 4.2 presents the ablation study on the detailed setting, especially the recursive depth of our CompConv. Sec. 4.3 and Sec. 4.4 evaluate the well designed CompConv on various visual benchmarks.

4.1. Experimental Settings

Dataset. ImageNet [5] contains around 1.28 million training images and 5k validation images. It includes 400 image categories in total. CIFAR-10 [20] consists of 50k training images and 10k testing images in 10 classes. The CIFAR-100 [20] increases the number of classes to 100, which is a finer version of CIFAR-10 [20]. The COCO benchmark [27] contains 118k images for training, 5k images for validation.

Training and Inference. For ImageNet classification, we adopt data augmentation scheme including random crop and mirroring following [11]. We use the momentum of 0.9 and synchronized SGD training over 8 GPUs. Each GPU has a batch-size of 32, resulting in a mini-batch of 256 in total. The learning rate is 0.1 and will be reduced by a factor of 10 at 30, 60, 90 epochs (100 epochs in total), respectively. At inference time, the shorter side of samples from ImageNet is resized to 256 and then cropped with 224×224 at the center region.

Table 1. **Effect of Shuffle Block.** We deploy CompConv on VGG-16 following Eq. (5). The proposed CompConv shuffle channels of feature maps after composition while the one without shuffle block keeps sequential concatenation.

Model	Params	FLOPs	Top-1(%)
w.o Shuffle-Block	3.3M	107M	92.8
Ours	3.3M	107M	93.8

For CIFAR-10 and CIFAR-100, we adopt simple data augmentation for training: 4 pixels are padded on each side, and a 32×32 crop is randomly sampled from the padded image or its horizontal flip. We start with a learning rate of 0.1 with a total batch size of 128, divide it by 20 at 60, 120 epochs (150 epochs in total). And we feed the single view of original images to the network at inference time.

For object detection on COCO [27], the input images are resized to a short side of 800 and a long side without exceeding 1333 following [3]. The models are initialized by the pre-trained ImageNet model and trained on 8 GPUs with 2 images per GPU for 12 epochs ($1 \times$ settings). In SGD training, the learning rate is initialized to 0.01 and then divided by 10 at epochs 8 and 11. The weight decay and momentum parameters are set to 10^{-4} and 0.9, respectively.

4.2. Ablation Study

We choose VGG16 [36] as the basic backbone network to conduct the following ablation experiments on CIFAR-10 [20]. All our experiments are conducted by replacing original convolutions with the same channels of input and output feature maps. And the new models are denoted as Comp-VGG16. Since VGG-16 is originally designed for ImageNet, we adopt a little modification to the structure like [29] which is widely used in the literature to handle samples from CIFAR-10 in lower resolution.

Effect of Shuffle Block. It’s easily noticed that the output feature maps are obtained by integrating all recursive results in CompConv. As shown in Eq. (3), the divided part of feature maps, \mathbf{X}_A , identically maps few channels of primary features and transmits them to compose the output of the module. To ensure \mathbf{X}_A does not include fixed feature maps with the constant position of channels, we adopt a shuffle block [42] in 4 groups after the concatenation. Tab. 1 compares the results between our CompConv and the one without shuffle block. The results imply that shuffle block can improve performance without adding any computations. The shuffle operation in CompConv prevents propagating the same feature from shallow feature to the final output and enhance the feature representation ability. Thus we’ll adopt the CompConv with shuffle block in the following experiments.

Effect of Identical Mapping. During recursive computation in CompConv, $\{\mathbf{X}_{A_i}, i \in \{1 \dots d\}\}$ is implemented

Table 2. **Effect of Identical Mapping.** We study the effectiveness of our *fixed* identity mapping. The results of *conv*, *random* and *group* are shown for comparison with our *fixed* identity mapping.

Model	Option	Params	FLOPs	Top-1(%)
Comp-VGG16	<i>group</i>	3.3M	107M	93.7
	<i>random</i>	3.3M	107M	92.7
	<i>conv</i>	3.6M	112M	93.8
Ours	<i>fixed</i>	3.3M	107M	93.8

by the identical mapping from the starting channels of input feature maps. We conduct the following ablation to study the effect of the identity mapping options for X_A . The basic operation to map starting channels from input feature map is denoted as *fixed* in Tab. 2.

We first give a comparison by generating these feature maps via applying 1×1 convolution to the input feature map. The *conv* option brings more computations while the performance with 93.16% is as same as the *fixed* mapping option with 93.17%.

Except for the *conv* option with larger parameters and FLOPs, we also give another two mapping options as group-selected and random-selected from the input feature without extra computation cost. The *group* option means dividing input channels by several groups then pick up specific feature maps from each group. With the same parameters and FLOPs, feature sampling from group-division shows similar classification accuracy with *fixed* option. Intuitively, we think that such channels assignment drives the CompConv to seek principle feature maps in the input regardless of positions. The *random* option implies no fixed assignment but in a random mode with slightly lower performance than other mapping options with 0.4% drop. Thus, we choose the simple way which incorporates the specific identical mapping from *fixed* starting position of the input feature map.

Analysis on Recursion Depth d . As shown in Eq. (8), our CompConv has one hyper-parameter d denotes recursion depth while larger d brings greater computation compression. To explore the effect of different d , we firstly adopt a global setting of different d by replacing all convolutions with CompConv. As shown in Tab. 4, setting with global $d = 1$ in CompConv can save half of parameters and FLOPs with 1% loss of Top-1 accuracy. When equipping CompConv with the global setting of $d = 2$ or $d = 3$, the whole network is compressed with higher efficiency but still maintains good performance. However, for $d = 4$ the performance presents a distinct drop compared with the other settings because of the severe squeeze for the primal channels based on Eq. (5). So we give the adaptive separation strategy shown in Sec. 3 on d , not exceeding 4 as written in Eq. (6).

Table 3. **Effect of Recursion Depth d .** We conduct experiments on global settings with different d in Comp-VGG16. SSAD denotes the Separation Strategy with Adaptive Depth.

Model	d	Params	FLOPs	Top-1(%)
VGG16	0	14.7M	314M	94.1
Comp-VGG16	1	7.4M	158M	93.0
	2	4.3M	100M	92.7
	3	2.9M	73M	92.6
	4	2.2M	56M	92.0
Ours(w. SSAD)	-	3.3M	107M	93.8

In addition, we substitute vanilla convolution at different stages to study the relation between d and channels. According to the four MaxPooling layers in VGG16, we divide the whole network into five stages. We deploy CompConv only at the first stage with 64 channels and last stage with 512 channels to conduct this comparative study. Comp-VGG16(1) indicates first stage replacement using CompConv by setting d into different values from $\{1, 2, 3\}$ while other regular convolutions keep unchanged. In this case, $d = 1, 2, 3$ achieves nearly the same compression performance, while the Top-1 accuracy for $d = 2, 3$ is lower than $d = 1$ with 0.7%, which supports that the former stages of the network are less redundant.

When we deploy CompConv with $d = 1, 2, 3$ only at the 5-th stage, it’s noticed that parameters and FLOPs of the network decrease greatly, which still maintains considerable accuracy yet. So we could use CompConv with a larger $d = 3$ at the high-dimensional stage to achieve a better trade-off between computing expense and accuracy. It reveals that there is more redundancy with larger channels. Consequently, following Eq. (6), the convolution with larger channels are equipped with larger recursive depth and vice versa. Based on an adaptive selection of d described in Sec. 3.4, our Comp-VGG16 demonstrates more significant performance than all with global d settings as well as saving a large amount of parameters and FLOPs. Thus we adopt the adaptive separation strategy for CompConv in the following experiments for effectiveness and efficiency.

4.3. Image Classification

After studying the superiority of the proposed CompConv module for efficient feature learning, we then evaluate the well-designed CompConv architecture on image classification, including various datasets including CIFAR-10, CIFAR-100, and ImageNet.

On CIFAR-10, we use the Comp-VGG16 model mentioned in Sec. 4.2 to compare with other competitors. For ResNet on CIFAR-100 and ImageNet, we conduct experiments by replacing regular convolutions with CompConv except for the **Conv1** before the first MaxPooling. As a

Table 4. **Effect of Recursion Depth d on Different Stages.** CompConv with different recursive depth d are adopted to the original VGG-16 at first stage and last stage of VGG-16 independently, named Comp-VGG16-1s and Comp-VGG16-5s. We use them to evaluate the effect of recursion depth d on different stages.

Model	d	Params	FLOPs	Top-1(%)
VGG16	0	14.7M	314M	94.1
Comp-VGG16-1s	1	14.7M	295M	93.8
	2	14.7M	289M	93.1
	3	14.7M	287M	93.1
Comp-VGG16-5s	1	8.2M	253M	93.8
	2	5.5M	228M	93.7
	3	4.3M	216M	93.6

summary of former configurations, we adopt the recommended separation strategy as Eq. (6). Unless specified otherwise, the training setting is as same as in Sec. 4.1. Other than the model in full precision, we also integrate CompConv with a quantized network [33] to show the generality of CompConv.

4.3.1 VGG on CIFAR-10

We compare our Comp-VGG16 with some state-of-the-art models, including different types of model compression approaches. i.e., l_1 pruning [22], channel pruning [12], HRank [23], SSS [17], SBP [12]. In comparison, our Comp-VGG16 can achieve the best accuracy (only drops 0.3% than the original one) with $3\times$ acceleration and can save nearly about $5\times$ storage. Specifically, HRank [23] can achieve the smallest number of parameters, while Comp-VGG16 saves 39M FLOPs with the 0.4% improvement on Top-1 accuracy.

4.3.2 Binarized CompConv on CIFAR-10

Except applying our CompConv to replace regular convolution in a full precision model, we also integrate CompConv with binary methods in XNORNet [33] to show the effectiveness of CompConv in the lightweight scenario. We perform the same binarization process as shown in [33] for vanilla VGG-16, Ghost-VGG16 and our Comp-VGG16 and make an apple-to-apple comparison with these methods. As shown in Tab. 6, the binarized CompConv beat binarized Ghost both in parameters and FLOPs, with nearly $2\times$ acceleration and storage saving. Though the binarized model is lightweight enough, there still exists redundancy of computation cost as pointed by [38], the computation efficiency of which can also be improved by replacing regular convolution with our CompConv.

Table 5. Comparison of state-of-the-art methods for compressing VGG-16 [36] on CIFAR-10 [20] dataset.

Model	Params	FLOPs	Top-1(%)
VGG16 [36]	15M	314M	94.1
L1-VGG16 [22]	5.4M	206M	93.4
SSS-VGG16 [17]	3.9M	183M	93.0
Ghost-VGG16 [9]	7.7M	158M	93.7
HRank-VGG16 [23]	2.5M	146M	93.4
SBP-VGG16 [22]	-	136M	92.5
Comp-VGG16	3.3M	107M	93.8

Table 6. Comparison with convolutional compressing design. We use another plug-and-play Ghost module [9] to make impartial comparison using the same binary method [33].

Model	Params	FLOPs	Top-1(%)
Bin-VGG16	1.3M	22.4M	83.8
Bin-Ghost-VGG16	0.9M	11.6M	79.1
Bin-Comp-VGG16	0.3M	6.6M	80.0

4.3.3 ResNet on CIFAR-100

We also apply CompConv on ResNet to verify its capability of finer classification on CIFAR-100. The samples of CIFAR-100 are as same as CIFAR-10 but with finer labels. Here we mainly compare with another similar lightweight convolutional design Ghost Module.

We replace CompConv in both BasicBlock and Bottleneck respectively in ResNet18 and ResNet50 illustrated in Fig. 3. According to results listed in Tab. 7, our CompConv achieves better accuracy (nearly 1%) but around 30% FLOPs reduction than GhostModule. The results listed in Tab. 5 demonstrates that there is a tiny gap between the original model and substitute one with CompConv. Intuitively, CIFAR-100 with finer labels is more difficult than CIFAR-10, which needs a larger model with more computation to maintain the feature representation. In spite of this, we still obtain a good result with 40% FLOPs and parameters of conventional ResNet at the cost of 3% performance drop.

4.3.4 ResNet on ImageNet

Finally, We conduct experiments for ResNet-50 on the challenging ImageNet dataset, comparing with other state-of-the-art methods. For this network structure, we replace all convolutions with CompConv as shown in Fig. 3, where we also take the example of separation strategy using Eq. (6) as CompConv128.

In order to give a comprehensive comparison on various FLOPs group, we apply three configurations by setting c_0 to 512, 256 and 128 in Eq. (6), which are named as Comp512-

Table 7. Comparison of Ghost-ResNet [9] on CIFAR-100 [20] dataset. We replace all convolutions in both basic block and bottleneck with Ghost module and CompConv within ResNet.

Model	Params	FLOPs	Top-1(%)
Res18	11.2M	0.6G	76.4
Ghost-Res18	5.7M	0.3G	72.3
Comp-Res18	3.3M	0.2G	73.1
Res50	23.7M	1.3G	77.4
Ghost-Res50	13.5M	0.7G	74.6
Comp-Res50	12.2M	0.5G	75.5

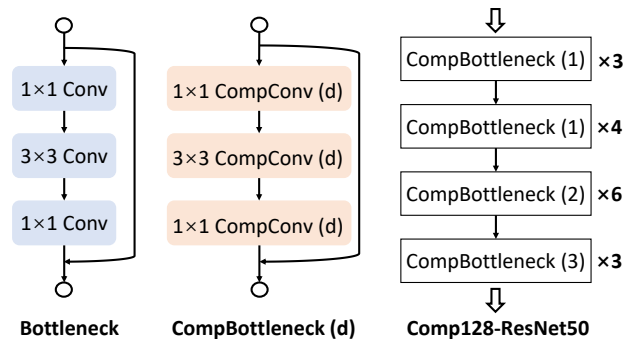


Figure 3. CompConv-ResNet50

ResNet50, Comp256-ResNet50, Comp128-ResNet50.

Compared with the recent state-of-the-art methods including Thinet[31], Versatile filters [37] and Sparse structure selection [17] in larger FLOPs group, our Comp512-ResNet50 beats them with 1.3%, 1.6%, 0.9% promotion of accuracy. Besides, our Comp512-ResNet50 keeps nearly the same performance as the original ResNet50 with $1.7\times$ acceleration. When we further set c_0 to 256, our Comp256-ResNet50 has a 0.6% accuracy drop while still obtains better results than HRank [23] and GhostModule [9] with nearly the same FLOPs and fewer parameters.

For the lower FLOPs group, our Comp128-ResNet50 achieves the highest accuracy with the lowest FLOPs and parameters. In contrast, compared methods including ABC-Pruner [24], GDP[25], GAL[26] with similar weights or FLOPs have much lower performance than ours.

All things considered, CompConv provides us with excellent compact convolutional design for CNN network.

Since we've talked thoroughly about the influence of recursion depth d when using CompConv, we adopt recommended separation strategy based on Eq. (6) with $c_0 = 512, 256, 128, 64, 32$ to get more results in Fig. 4. This solution gives a dynamic comparison by setting CompConv-ResNet50 with different computation-consuming levels. For comparison, we select several state-of-the-art methods to draw the FLOPs v.s. Top-1 accuracy curve. From an

Table 8. Comparison of state-of-the-art methods for compressing ResNet50 on ImageNet.

Model	Params	FLOPs	Top-1(%)	Top-5(%)
ResNet50 [11]	25.6M	4.1G	76.1	92.9
Versatile-ResNet50 [37]	11.0M	3.0G	74.5	91.8
SSS-ResNet50-32 [17]	18.6M	2.8G	74.2	91.9
Thinet-ResNet50 [31]	16.9M	2.6G	72.1	90.3
Comp512-ResNet50	15.3M	2.4G	75.8	92.8
HRank-ResNet50 [23]	16.2M	2.3G	75.0	92.3
NISP-ResNet50 [40]	14.4M	2.3G	-	90.8
Ghost-ResNet50-s2 [9]	13.0M	2.2G	75.0	92.3
Comp256-ResNet50	13.7M	2.1G	75.2	92.3
GDP-ResNet50-0.6 [25]	-	1.9G	71.2	90.7
GAL-ResNet50 [26]	19.3M	1.8G	71.8	90.8
ABC-ResNet50 [24]	11.2M	1.8G	73.5	91.5
HRank-ResNet50 [23]	13.8M	1.6G	72.0	91.0
Comp128-ResNet50	8.7M	1.6G	73.7	91.2

Table 9. Results on COCO with Faster-RCNN detector.

Model	Params	FLOPs	mAP
ResNet50	41.5M	207.1G	37.4
Comp512-ResNet50	31.5M	176.3G	37.2
Comp128-ResNet50	30.1M	160.8G	36.6

overall view of Fig. 4, the curve of our CompConv are above all other methods including SSS[17], Thinet[31], GDP[25] and GAL[26], ABC-Prunner [24], HRANK [23]. It shows that our CompConv gains higher accuracy with the same computation cost. As for the same performance, the model with our CompConv performs more compact than others.

4.4. Object Detection

In addition to results on classification benchmarks, we also evaluate our CompConv in object detection on MS COCO dataset [27]. We deploy our CompConv to the state-of-the-art detector Faster-RCNN [34] following the train and inference setting in Sec. 3.4. As shown in Tab. 9 where the FLOPs are calculated using 800×1333 images following MMDetection [3], our Comp512-ResNet50 achieve nearly the same performance with lower computational costs compared with vanilla ResNet50. The smaller model Comp128-ResNet50 can save 47G FLOPs and 11.43M parameters cost of a slight drop on mAP.

5. Discussion and Conclusion

To reduce the computational costs of CNN, this paper presents a novel design denoted as CompConv from perspective of the basic convolution unit. It utilizes divide-and-conquer strategy to simplify transformation of feature maps. The proposed CompConv is a plug-and-play module for re-

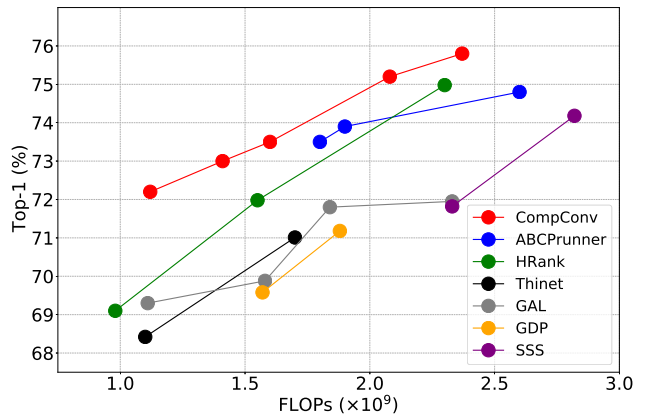


Figure 4. Top1-accuracy v.s. FLOPs for ResNet-50 Model on ImageNet.

placing vanilla convolution without additional restrictions. Comprehensive experiments conducted on various visual benchmark models illustrate that our CompConv greatly save FLOPs and parameters while still keeping a considerable accuracy performance with good transferability. It's an efficient convolutional design to make CNNs lighter to deploy on resources-limited ends.

One important future work direction is to design a more adaptive and comprehensive solution when applying CompConv to wider network. The recursion depth d accounts a lot in saving computation and enhancing performance. Maybe we can use AutoML technology to always find out the most suitable settings when applying to different networks. In that way, CompConv would play a stable role in any convolutional network structures more easily. Besides, how to deploy the CompConv in a smaller model *i.e.* the pruned model or the quantized model to further compress the network remains an open problem.

References

- [1] Gabriel Bender, Pieter-Jan Kindermans, Barret Zoph, Vijay Vasudevan, and Quoc Le. Understanding and simplifying one-shot architecture search. In *ICML*, 2018. 2
- [2] Han Cai, Chuang Gan, Tianzhe Wang, Zhekai Zhang, and Song Han. Once-for-all: Train one network and specialize it for efficient deployment. *arXiv preprint arXiv:1908.09791*, 2019. 2
- [3] Kai Chen, Jiaqi Wang, Jiangmiao Pang, Yuhang Cao, Yu Xiong, Xiaoxiao Li, Shuyang Sun, Wansen Feng, Ziwei Liu, Jiarui Xu, Zheng Zhang, Dazhi Cheng, Chenchen Zhu, Tianheng Cheng, Qijie Zhao, Buyu Li, Xin Lu, Rui Zhu, Yue Wu, Jifeng Dai, Jingdong Wang, Jianping Shi, Wanli Ouyang, Chen Change Loy, and Dahua Lin. MMDetection: Open mmlab detection toolbox and benchmark. *arXiv preprint arXiv:1906.07155*, 2019. 5, 8
- [4] Matthieu Courbariaux, Yoshua Bengio, and Jean-Pierre David. Binaryconnect: Training deep neural networks with binary weights during propagations. In *Adv. Neural Inform. Process. Syst.*, 2015. 2
- [5] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2009. 4
- [6] Xin Dong, Shangyu Chen, and Sinno Pan. Learning to prune deep neural networks via layer-wise optimal brain surgeon. In *Adv. Neural Inform. Process. Syst.*, 2017. 2
- [7] Suyog Gupta, Ankur Agrawal, Kailash Gopalakrishnan, and Pritish Narayanan. Deep learning with limited numerical precision. In *ICML*, 2015. 2
- [8] Bo Han, Quanming Yao, Xingrui Yu, Gang Niu, Miao Xu, Weihua Hu, Ivor Tsang, and Masashi Sugiyama. Co-teaching: Robust training of deep neural networks with extremely noisy labels. In *Adv. Neural Inform. Process. Syst.*, 2018. 1, 2
- [9] Kai Han, Yunhe Wang, Qi Tian, Jianyuan Guo, Chunjing Xu, and Chang Xu. Ghostnet: More features from cheap operations. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2020. 1, 2, 7, 8
- [10] Song Han, Huizi Mao, and William J Dally. Deep compression: Compressing deep neural networks with pruning, trained quantization and Huffman coding. *Int. Conf. Learn. Represent.*, 2015. 1, 2
- [11] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2016. 1, 4, 8
- [12] Yihui He, Xiangyu Zhang, and Jian Sun. Channel pruning for accelerating very deep neural networks. In *Int. Conf. Comput. Vis.*, 2017. 1, 2, 6
- [13] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015. 1, 2
- [14] Andrew Howard, Mark Sandler, Grace Chu, Liang-Chieh Chen, Bo Chen, Mingxing Tan, Weijun Wang, Yukun Zhu, Ruoming Pang, Vijay Vasudevan, et al. Searching for mobilenetv3. In *Int. Conf. Comput. Vis.*, 2019. 2
- [15] 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 preprint arXiv:1704.04861*, 2017. 1, 2
- [16] Zehao Huang and Naiyan Wang. Like what you like: Knowledge distill via neuron selectivity transfer. *arXiv preprint arXiv:1707.01219*, 2017. 2
- [17] Zehao Huang and Naiyan Wang. Data-driven sparse structure selection for deep neural networks. In *Eur. Conf. Comput. Vis.*, 2018. 2, 6, 7, 8
- [18] Forrest N Iandola, Song Han, Matthew W Moskewicz, Khalid Ashraf, William J Dally, and Kurt Keutzer. Squeezenet: Alexnet-level accuracy with 50x fewer parameters and 0.5 mb model size. *arXiv preprint arXiv:1602.07360*, 2016. 1, 2
- [19] Benoit Jacob, Skirmantas Kligys, Bo Chen, Menglong Zhu, Matthew Tang, Andrew Howard, Hartwig Adam, and Dmitry Kalenichenko. Quantization and training of neural networks for efficient integer-arithmetic-only inference. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2018. 2
- [20] Alex Krizhevsky et al. Learning multiple layers of features from tiny images. 2009. 4, 5, 7
- [21] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Communications of the ACM*, 2012. 1
- [22] Hao Li, Asim Kadav, Igor Durdanovic, Hanan Samet, and Hans Peter Graf. Pruning filters for efficient convnets. *Int. Conf. Learn. Represent.*, 2017. 1, 2, 6, 7
- [23] Mingbao Lin, Rongrong Ji, Yan Wang, Yichen Zhang, Baochang Zhang, Yonghong Tian, and Ling Shao. Hrank: Filter pruning using high-rank feature map. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2020. 2, 6, 7, 8
- [24] Mingbao Lin, Rongrong Ji, Yuxin Zhang, Baochang Zhang, Yongjian Wu, and Yonghong Tian. Channel pruning via automatic structure search. *arXiv preprint arXiv:2001.08565*, 2020. 2, 7, 8
- [25] Shaohui Lin, Rongrong Ji, Yuchao Li, Yongjian Wu, Feiyue Huang, and Baochang Zhang. Accelerating convolutional networks via global & dynamic filter pruning. In *IJCAI*, 2018. 2, 7, 8
- [26] Shaohui Lin, Rongrong Ji, Chenqian Yan, Baochang Zhang, Liujuan Cao, Qixiang Ye, Feiyue Huang, and David Doermann. Towards optimal structured CNN pruning via generative adversarial learning. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2019. 2, 7, 8
- [27] 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 *Eur. Conf. Comput. Vis.*, 2014. 4, 5, 8
- [28] Chenxi Liu, Barret Zoph, Maxim Neumann, Jonathon Shlens, Wei Hua, Li-Jia Li, Li Fei-Fei, Alan Yuille, Jonathan Huang, and Kevin Murphy. Progressive neural architecture search. In *Eur. Conf. Comput. Vis.*, 2018. 2
- [29] Kuang Liu. Pytorch cifar10. <https://github.com/kuangliu/pytorch-cifar.git>, 2019. 5

- [30] Zhuang Liu, Jianguo Li, Zhiqiang Shen, Gao Huang, Shoumeng Yan, and Changshui Zhang. Learning efficient convolutional networks through network slimming. In *Int. Conf. Comput. Vis.*, 2017. 2
- [31] Jian-Hao Luo, Jianxin Wu, and Weiyao Lin. Thinet: A filter level pruning method for deep neural network compression. In *Int. Conf. Comput. Vis.*, 2017. 2, 7, 8
- [32] Ningning Ma, Xiangyu Zhang, Hai-Tao Zheng, and Jian Sun. Shufflenet v2: Practical guidelines for efficient cnn architecture design. In *Eur. Conf. Comput. Vis.*, 2018. 2
- [33] Mohammad Rastegari, Vicente Ordonez, Joseph Redmon, and Ali Farhadi. Xnor-net: Imagenet classification using binary convolutional neural networks. In *Eur. Conf. Comput. Vis.*, 2016. 2, 6, 7
- [34] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Adv. Neural Inform. Process. Syst.*, 2015. 8
- [35] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2018. 1, 2
- [36] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 1, 4, 5, 7
- [37] Yunhe Wang, Chang Xu, XU Chunjing, Chao Xu, and Dacheng Tao. Learning versatile filters for efficient convolutional neural networks. In *Adv. Neural Inform. Process. Syst.*, 2018. 7, 8
- [38] Yinghao Xu, Xin Dong, Yudian Li, and Hao Su. A main/subsidiary network framework for simplifying binary neural networks. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2019. 6
- [39] Junho Yim, Donggyu Joo, Jihoon Bae, and Junmo Kim. A gift from knowledge distillation: Fast optimization, network minimization and transfer learning. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2017. 2
- [40] Ruichi Yu, Ang Li, Chun-Fu Chen, Jui-Hsin Lai, Vlad I Morariu, Xintong Han, Mingfei Gao, Ching-Yung Lin, and Larry S Davis. Nisp: Pruning networks using neuron importance score propagation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2018. 2, 8
- [41] Sergey Zagoruyko and Nikos Komodakis. Paying more attention to attention: Improving the performance of convolutional neural networks via attention transfer. *arXiv preprint arXiv:1612.03928*, 2016. 2
- [42] Xiangyu Zhang, Xinyu Zhou, Mengxiao Lin, and Jian Sun. Shufflenet: An extremely efficient convolutional neural network for mobile devices. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2018. 2, 5
- [43] Zhao Zhong, Junjie Yan, Wei Wu, Jing Shao, and Cheng-Lin Liu. Practical block-wise neural network architecture generation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2018. 2
- [44] Aojun Zhou, Anbang Yao, Yiwen Guo, Lin Xu, and Yurong Chen. Incremental network quantization: Towards lossless cnns with low-precision weights. *arXiv preprint arXiv:1702.03044*, 2017. 2
- [45] Barret Zoph and Quoc V Le. Neural architecture search with reinforcement learning. *arXiv preprint arXiv:1611.01578*, 2016. 2
- [46] Barret Zoph, Vijay Vasudevan, Jonathon Shlens, and Quoc V Le. Learning transferable architectures for scalable image recognition. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2018. 2