

# GroSS: Group-Size Series Decomposition for Grouped Architecture Search

Henry Howard-Jenkins<sup>[0000-0003-3914-5883]</sup>, Yiwen Li<sup>[0000-0002-7794-9391]</sup>, and  
Victor Adrian Prisacariu<sup>[0000-0002-0630-6129]</sup>

Active Vision Laboratory, University of Oxford, UK  
{henryhj, kate, victor}@robots.ox.ac.uk

**Abstract.** We present a novel approach which is able to explore the configuration of grouped convolutions within neural networks. Group-size Series (GroSS) decomposition is a mathematical formulation of tensor factorisation into a series of approximations of increasing rank terms. GroSS allows for dynamic and differentiable selection of factorisation rank, which is analogous to a grouped convolution. Therefore, to the best of our knowledge, GroSS is the first method to enable simultaneous training of differing numbers of groups within a single layer, as well as all possible combinations between layers. In doing so, GroSS is able to train an entire grouped convolution architecture search-space concurrently. We demonstrate this through architecture searches with performance objectives on multiple datasets and networks. GroSS enables more effective and efficient search for grouped convolutional architectures.

**Keywords:** Group Convolution, Network Acceleration, Architecture Search

## 1 Introduction

In recent years, there has been a flurry of deep neural networks (DNNs) producing remarkable results on a broad variety of tasks. In particular, grouped convolution has become a widely used tool in some prevalent networks. ResNeXt [28] used grouped convolution for improved accuracy over the analogous ResNets [7]. On the other hand, Xception [3], MobileNet [8], and various others [29, 19] have used depthwise convolutions, which are the special case of grouped convolutions where the number of groups is equal to the number of in channels, in a for extremely low-cost inference. With these architectures, grouped convolution has proven to be a valuable design tool for high-performance and low-cost design alike. But, its application to these contrasting performance profiles has so far, to the best of our knowledge, remained relatively unexplored.

Finding a heuristic or intuition for how combinations of grouped convolutions with varying numbers of groups interact within a network is challenging. Grouped convolution, therefore, presents itself as an ideal candidate for Neural Architecture Search (NAS), which has provided an alternative to hand designed networks. NAS allows for the search and even direct optimisation of the network’s structure. But, the search space for architectures is often vast, with potentially limitless design choices. Furthermore, each configuration must undergo

some training or fine-tuning for its efficacy to be determined. This has led to the development of methods which lump multiple design parameters together, which reduce the search space in a principled manner [23], as well as creating the need for sophisticated search algorithms [17, 27], which can more quickly converge to an improved design. Both techniques reduce the number of search iterations and ultimately reduce the number of required training/fine-tuning stages.

In this work, however, we do not wish to make assumptions about the grouped convolution manifold. We achieve this with the introduction of a Group-size Series (GroSS) decomposition. GroSS allows us to train the entire search space of architectures *simultaneously*. In doing so, we shift the expense of architecture search with respect to groups away from decomposition and training, and towards cheaper test-time sampling. This allows for the exploration of possible configurations, while significantly reducing the need for imparting bias on the group design hyperparameter selection.

The contributions of this paper can be summarised as follows:

1. We present GroSS decomposition – a novel formulation of tensor decomposition as a series of rank approximations. This provides a mathematical basis for grouped convolution as a series of increasing rank terms.
2. GroSS provides the apparatus for differentially switching between grouped convolution ranks. Therefore, to the best of our knowledge, it is the first simultaneous training of differing numbers of groups within a single layer, as well as the all possible configurations between layers. This makes feasible, for the first time, a search for rank selection for network compression.
3. We explore this concurrently-trained architecture space in the context of network acceleration. We factorise a small network, as well VGG-16 and ResNet-18, and propose exhaustive and breadth-first searches on CIFAR-10 and ImageNet. We demonstrate the efficacy of the GroSS for rank selection search over a more conventional approach of partial training schedules.

## 2 Related Work

Grouped convolution has had a wide impact on neural network architectures, particularly due to its efficiency. It was first introduced in AlexNet [14] as an aid for the single network to be trained over multiple GPUs. Since then, it has had a wide impact on DNN architecture design. Deep Roots [9] was the first to introduce group convolution for efficiency, while ResNeXt [28] used grouped convolutions synonymously with concept of *cardinality*, ultimately exploiting the efficiency of grouped convolutions for high-accuracy network design. The reduced complexity of grouped convolution allowed for ResNeXt to incorporate deeper layers within the ResNet-analogous residual blocks [7]. In all, this allowed higher accuracy with a similar inference cost as an equivalent ResNet. The efficiency of grouped convolution has also led to several low-cost network designs. Sifre [20] first introduced depthwise separable convolutions, which were later utilised by Xception [3]. MobileNet [8] utilised a ResNet-like bottleneck design with depthwise convolutions for an extremely efficient network with mobile applications

in mind. ShuffleNet [29] was also based on a depthwise bottleneck, however, pointwise layers were also made grouped convolutions.

Previous works [10, 5, 15, 25] have applied low-rank approximation of convolution for network compression and acceleration. Block Term Decomposition (BTD) [4] has recently been applied to the task of network factorisation [2], where it was shown that the BTD factorisation of a convolutional weight was equivalent to a grouped convolution within a bottleneck architecture. Wang *et al.* [26] applied this equivalency for network acceleration. Since decomposition is costly, these methods have relied on heuristics and intuition to set hyperparameters such as the rank of successive layers within the decomposition. In this paper, we present a method for decomposition which allows for exploration of the decomposition hyperparameters and all the combinations.

Existing architecture search methods have overwhelmingly favoured reinforcement learning. Examples of this include, but are not limited to, NAS-Net [30], MNasNet [22], ReLeq-Net [6]. In broad terms, these methods all set a baseline structure, which is manipulated by a separate controller. The controller optimises the structure through an objective based on network performance. There has also been work in differentiable architecture search [27, 17] which makes the network architecture manipulations themselves differentiable. In addition, work such as [23] aims to limit the network scaling within a performance envelope to a single parameter.

These methods all have a commonality: the cost of re-training or fine-tuning at each stage motivates the recovery of the optimal architecture in as few training steps as possible, whether this is achieved through a trained controller, direct optimisation or significantly reducing the search space. In this work, however, GroSS allows efficient weight-sharing between varying grouped architectures, thus enabling them to be trained at once. This is similar to the task of one-shot architecture search. SMASH [1] use a hypernetwork to predict weights for each architecture. The work of Li *et al.* [16] bears most resemblance to this paper, where randomly sampled architectures are used to train shared-weights.

### 3 Method

In this section, we will first introduce Block Term Decomposition (BTD) and detail how its factorisation can be applied to a convolutional layer. After that, we will introduce GroSS decomposition, where we formulate a unification of a series of ranked decompositions so that they can dynamically and differentially be combined. We detail the training strategy for training the whole series at once. Finally, we detail the methodology of our exhaustive and breadth-first search.

#### 3.1 General Block Term Decomposition

Block Term Decomposition (BTD) [4] aims to factorise a tensor into the sum of multiple low rank-Tuckers [24]. That is, given an  $N^{\text{th}}$  order tensor  $\mathbf{X} \in$

$\mathbb{R}^{d_1 \times d_2 \times \dots \times d_N}$ , BTD factorises  $\mathbf{X}$  into the sum of  $R$  terms with rank  $(d'_1, d'_2, \dots, d'_N)$ :

$$\mathbf{X} = \sum_{r=1}^R \mathbf{G}_r \times_1 \mathbf{A}_r^{(1)} \times_2 \mathbf{A}_r^{(2)} \times_3 \dots \times_N \mathbf{A}_r^{(N)} \quad (1)$$

where  $\left\{ \begin{array}{l} \mathbf{G} \in \mathbb{R}^{d'_1 \times d'_2 \times \dots \times d'_N} \\ \mathbf{A}_r^{(n)} \in \mathbb{R}^{d_n \times d'_n}, n \in \{1, \dots, N\} \end{array} \right.$

In the above,  $\mathbf{G}$  is known as the *core* tensor and we will refer to  $\mathbf{A}$  as *factors* matrices. We use the usual notation  $\times_n$  to represent the *mode- $n$*  product [4].

### 3.2 Converting a Single Convolution to a Bottleneck Using BTD

Here, we can restrict discussion from a general,  $N$ -mode, tensor to the 4-mode weights of a 2D convolution as follows:  $\mathbf{X} \in \mathbb{R}^{t \times u \times v \times w}$ , where  $t$  and  $u$  represent the number of input and output channels, and  $v$  and  $w$  the spatial size of the filter kernel. Typically the spatial extent of each filter is small and thus we only factorise  $t$  and  $u$ . To eliminate superscripts, we define  $\mathbf{B} = \mathbf{A}^{(1)}$  and  $\mathbf{C} = \mathbf{A}^{(2)}$ . Therefore, the BTD for convolutional weights is expressed as follows:

$$\mathbf{X} = \sum_{r=1}^R \mathbf{G}_r \times_1 \mathbf{B}_r \times_2 \mathbf{C}_r \quad (2)$$

where  $\left\{ \begin{array}{l} \mathbf{G} \in \mathbb{R}^{t' \times u' \times v \times w} \\ \mathbf{B} \in \mathbb{R}^{t \times t'} \\ \mathbf{C} \in \mathbb{R}^{u \times u'} \end{array} \right.$

This factorisation of the convolutional weights into  $R$  groups forms a three-layer bottleneck-style structure [2]: a pointwise  $(1 \times 1)$  convolution  $\mathbf{P} \in \mathbb{R}^{t \times (Rt') \times 1 \times 1}$ , formed from factor  $\mathbf{B}$ ; followed by a grouped convolution  $\mathbf{R} \in \mathbb{R}^{t' \times (Ru') \times v \times w}$ , formed from core  $\mathbf{G}$  and with  $R$  groups; and finally another pointwise convolution  $\mathbf{Q} \in \mathbb{R}^{(Ru') \times u \times 1 \times 1}$ , formed from factor  $\mathbf{C}$ . With careful selection of the BTD parameters, the bottleneck approximation can be applied to any standard convolutional layer. This is visualised in Figure 1.

In Table 1, we detail how the dimensions of the bottleneck architecture are determined from its corresponding convolutional layer, and indicate how properties such as stride, padding and bias are applied within the bottleneck for

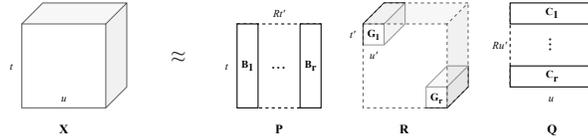


Fig. 1: Formation of bottleneck layers  $\mathbf{P}$ ,  $\mathbf{R}$  and  $\mathbf{Q}$  from BTD cores and factors.

Table 1: Convolution to grouped bottleneck. The table states how the convolutional parameters are used in the equivalent bottleneck

	Filter Size	$C_{in}$	$C_{out}$	Groups	Bias	Stride	Padding
Convolution	$v \times w$	$t$	$u$	1	$B$	$S$	$P$
	$1 \times 1$	$t$	$Rt'$	1	-	1	0
Bottleneck	$v \times w$	$Rt'$	$Ru'$	$R$	-	$S$	$P$
	$1 \times 1$	$Ru'$	$u$	1	$B$	1	0

equivalency with the original layer. It is worth noting that we often refer to the quantities  $t'$  or  $u'$  as the group-size; this quantity determines the number of channels in each group and is equivalent to the rank of the decomposition.

### 3.3 Group-size Series Decomposition

Group-size Series (GroSS) decomposition unifies multiple ranks of BTD factorisations. This is achieved by defining each successive factorisation relative to the lower order ranks. Thus we ensure that higher rank decompositions only contain information that was missed by the lower order approximations. Therefore the  $i^{\text{th}}$  approximation of  $\mathbf{X}$  is given as follows:

$$\mathbf{X} = \sum_{r=1}^{R_i} [(\mathbf{g}_r)_i + (\mathbf{G}'_r)_{i-1}] \times_1 [(\mathbf{b}_r)_i + (\mathbf{B}'_r)_{i-1}] \times_2 [(\mathbf{c}_r)_i + (\mathbf{C}'_r)_{i-1}]$$

$$\text{where } \begin{cases} (\mathbf{g}_r)_i, (\mathbf{G}'_r)_{i-1} \in \mathbb{R}^{t'_i \times u'_i \times v \times w} \\ (\mathbf{b}_r)_i, (\mathbf{B}'_r)_{i-1} \in \mathbb{R}^{t \times t'_i} \\ (\mathbf{c}_r)_i, (\mathbf{C}'_r)_{i-1} \in \mathbb{R}^{u \times u'_i} \end{cases} \quad (3)$$

Omitting  $r$  from the notation,  $\mathbf{g}_i$ ,  $\mathbf{b}_i$  and  $\mathbf{c}_i$  represent the additional information captured between the  $(i-1)^{\text{th}}$  and  $i^{\text{th}}$  rank of approximation, and  $\mathbf{G}'_{(i-1)}$ ,  $\mathbf{B}'_{(i-1)}$  and  $\mathbf{C}'_{(i-1)}$  to represent total approximation from lower rank approximations in the form of cores and factors. However, both the core and factors must be recomputed so that the dimensions match the ranks required  $R_i$ , which is not a trivial manipulation.

Instead, we introduce a function,  $\Psi_{g \rightarrow h}()$ , which allows the weights of a grouped convolution to be “expanded”. The expanded weight from a convolution with group-size  $g$  can be used in a convolution with group-size  $h$ , where  $h > g$ , giving identical outputs:

$$\mathbf{W}_g *_{g} \mathbf{F} \equiv \Psi_{g \rightarrow h}(\mathbf{W}_g) *_{h} \mathbf{F} \quad (4)$$

where  $\mathbf{W}_g$  is the weight for a grouped convolution,  $*_g$  refers to convolution with group-size  $g$ , and  $\mathbf{F}$  is the feature map to which the convolution is applied. We provide an example visualisation of  $\Psi()$  in Figure 2. This allows us to conveniently reformulate the GroSS decomposition in terms of the successive convolutional weights obtained from BTD, rather than within the cores and factors

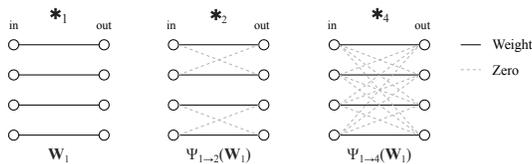


Fig. 2: Visualisation of a depthwise weight,  $\mathbf{W}_1$ , expanded for convolutions with groups of size 2 and 4.

directly. More specifically, we define the bottleneck weights for the  $N^{\text{th}}$  order GroSS decomposition with group-sizes,  $S = \{s_1, \dots, s_N\}$ , as follows:

$$\begin{aligned} \mathbf{R}_N &= \Psi_{s_1 \rightarrow s_N}(\mathbf{R}_1) + \sum_{i=2}^N \Psi_{s_i \rightarrow s_N}(\mathbf{r}_i) \\ \mathbf{P}_N &= \mathbf{P}_1 + \sum_{i=2}^N \mathbf{p}_i, \quad \mathbf{Q}_N = \mathbf{Q}_1 + \sum_{i=2}^N \mathbf{q}_i \end{aligned} \quad (5)$$

$\mathbf{R}_1$ ,  $\mathbf{P}_1$  and  $\mathbf{Q}_1$  represent the weights obtained from the lowest rank decomposition present in the series.  $\mathbf{r}_i$ ,  $\mathbf{p}_i$  and  $\mathbf{q}_i$  represent the additional information that the  $i^{\text{th}}$  rank decomposition contribute to the bottleneck approximation:

$$\mathbf{p}_i = \mathbf{P}_i - \mathbf{P}_{(i-1)}, \quad \mathbf{r}_i = \mathbf{R}_i - \Psi_{s_{(i-1)} \rightarrow s_i}(\mathbf{R}_{(i-1)}), \quad \mathbf{q}_i = \mathbf{Q}_i - \mathbf{Q}_{(i-1)}. \quad (6)$$

This formulation involving only manipulation of the convolutional weights is exactly equivalent to forming the bottleneck components  $\mathbf{r}_i$ ,  $\mathbf{p}_i$  and  $\mathbf{q}_i$  from  $\mathbf{g}_i$ ,  $\mathbf{b}_i$  and  $\mathbf{c}_i$ , as in the general BTD to bottleneck case.

Further, the grouped convolution weight expansion,  $\Psi()$ , enables us to dynamically, and differentially, change the group-size of a convolution. In itself, this is not particularly useful: a convolution with a larger group-size is requires more operations and more memory, while yielding identical outputs. However, it allows for direct interaction between differently ranked network decomposition and, therefore, the representation of one rank by the combination of lower ranks. Thus, GroSS treats the decomposition of the original convolution as the sum of successive order approximations, with each order contributing additional representational power.

**Training GroSS Simultaneously** The expression of a group-size  $s_i$  decomposition as the combination of lower rank decompositions is useful because it enables the group-size to be dynamically changed during training. The expansion and summation of convolutional weights is differentiable and so training at a high rank, also optimises the lower rank approximations simultaneously. To the best of our knowledge GroSS is the first method that allows weight-sharing between, and training of, convolutions with varying numbers of groups.

We leverage the series form of the factorisation during training, by randomly sampling a group-size for each decomposed layer at each iteration. We sample a group-size  $s_i$  for each decomposed layer uniformly. Through uniform sampling, we are able to train each network configuration equally.

### 3.4 Search

The objective for all the searches performed within in this paper is: given a base configuration, we aim to find an alternative configuration which is more accurate, but offers the same or cheaper inference. We implement two forms of search to achieve this leveraging GroSS decomposition: exhaustive and breadth-first.

**Exhaustive.** Within the exhaustive search, all possible configurations are evaluated. In this search, we simply filter any configuration with multiply accumulates (MACs) above the respective base configuration. After filtering, we can select the highest accuracy remaining.

**Breadth-first Search.** Where an exhaustive search is not feasible due to the sheer number of possible configurations, we use a greedy breadth-first search. We first randomly select a configuration which requires fewer operations for inference than the base configuration. We evaluate all neighbouring configurations—those which only require one layer to have its group-size changed—of the currently selected configuration. We select the neighbour with the highest accuracy that does not exceed the number of MACs as the base configuration for the next step. We repeat this step for a maximum of 25 times, or until there are no more accurate neighbours not exceeding the cost of the base configuration.

This is repeated 20 times. The most accurate configuration from all of the 20 runs is considered the result of the search. Since the search for a base configuration with fewer MACs is contained within the search-space with a higher limit, we perform them incrementally. This results in the same search process, but the first 10 runs are initialised using the top-10 highest accuracy results from the smaller search. We found that this generally led to faster stopping.

## 4 Application of GroSS

In this section, we explain how we apply GroSS to a several models across datasets. We first detail the dataset on which evaluation is conducted. Next, we describe the network architecture on which perform GroSS decomposition. Finally, we list the procedure for the decomposition and fine-tuning.

### 4.1 Datasets

We perform our experimental evaluation on CIFAR-10 [13] and ImageNet [14]. CIFAR-10 is a dataset consisting of 10 classes. The size of each image is  $32 \times 32$ . In

total there are 60,000 images, which are split into 50,000 train images and 10,000 testing images. We further divide the training set into a training and validation splits with 40,000 and 10,000 images, respectively. ImageNet consists of 1000 classes, with 1.2 million training images and 50,000 validation images. Since the test annotations are not available, we report our accuracy on the validation set.

## 4.2 Models

In this paper, we perform on three general network architectures: a custom 4-layer network, VGG-16 [21], and ResNet-18 [7]. Here, we provide an overview of the network definitions, with more details in the supplementary material.

Our 4-layer network has four convolutional layers, with output channel dimensions of 32, 32, 64 and 64, followed by two fully-connected layers of size 256 and 10. In our ImageNet experiments, we use a standard VGG-16 and ResNet-18, identical to those in [21] and [7], respectively. However, we make some changes to the fully connected structure in VGG-16 for training and inference on CIFAR-10. The convolutional layers instead followed by a  $2 \times 2$  max-pooling and two fully-connected layers of size 512 and 10, respectively. A ReLU layer and dropout with probability of 0.5 is applied between the fully-connected layers.

## 4.3 Decomposition

We perform GroSS decomposition on our small 4-layer network, as well as VGG-16 [21]. In each case we decompose all convolutional layers in the network aside from the first. Unless otherwise stated, we set the bottleneck width equal to the number of input channels. For the 4-layer network, group-sizes are set to all powers of 2 which do not exceed the bottleneck width for that respective layer. This leads to a total of 252 configurations represented by our decomposition. We decompose each layer in VGG-16 and ResNet-18 into 4 group-sizes: (1, 4, 16, 32). This leads to a total of  $4^{12}$  and  $4^{16}$  configurations represented by the decomposed VGG-16 and ResNet-18, respectively.

Our formulation of GroSS decomposition as a series of convolutional weight differences (expanded weights in the case of the grouped convolution), as detailed by Equation 5 means that we are able to use an off-the-shelf BTM framework [12]. For each group-size, we set the stopping criteria for BTM identically: when the decrease in approximation error between steps is below  $1 \times 10^{-6}$  for the 4-layer network and  $1 \times 10^{-5}$  for VGG-16, or  $5 \times 10^5$  steps have elapsed. We define approximation error as the Frobenius norm between the original tensor and the product of the BTM cores and factors divided by the Frobenius norm of the original tensor. For the 4-layer network, we perform this decomposition 5 times.

## 4.4 Fine-tuning

**CIFAR-10.** After we have performed GroSS decomposition on the network, we then fine-tune on the classification task. For the 4-layer network, we tune

for 150 epochs with a batch-size of 256, an initial learning rate of 0.0001 and momentum 0.9. We decay the learning rate by a factor of 0.1 after both 80 and 120 epochs. For VGG-16, we fine-tune with the same SGD parameters and batch-size, however we train for 200 epochs, and decay the learning rate after 100 and 150 epochs. All network parameters are frozen aside from the GroSS decomposition weights. During training, there is a 0.5 probability of horizontal flipping, zero-padding of size 2 is applied around all borders and a random  $32 \times 32$  crop is taken from the resulting image.

**ImageNet.** We decompose VGG-16 and ResNet-18 before funetuning on ImageNet. For VGG-16, we train using SGD for a total of 4 epochs with a batch-size of 128, leading to approximately  $10^4$  iterations. The initial learning rate is set to  $10^{-5}$ , which is decayed by a factor of 0.1 after 2 epochs. Momentum is set to 0.9. Again, all the network parameters are frozen, aside from the decomposition weights. For ResNet-18, we train for 8 epochs in total with a batch size of 512. The initial learning rate is set to  $5 \times 10^{-5}$ , with decay every 2 epochs. The images are resized so that the smallest side is of size 256. During training, the resized images are flipped horizontally with a probability of 0.5 and a random  $224 \times 224$  crop is taken. During testing, we simply take a centre crop from the resized image, hence evaluating 1-crop accuracy.

**Individual Configurations.** For the decomposition in the conventional manner, *i.e.* a singular group-size configuration, we decompose using exactly the same routine as with GroSS. However, the fine-tuning schedules are slightly modified. On CIFAR-10, we reduce the schedule for our 4-layer network and our CIFAR VGG-16 to 100 epochs. The initial learning rate is increased to 0.001, and decayed at 80 epochs. On Imagenet, we do not freeze the non-decomposed layers. The VGG-16 configurations have a schedule of 6 epochs, with learning rate decay occurring after every 2 epochs. The initial learning rate is kept at  $10^{-4}$ . For ResNet-18 configurations, we increase batch size to 512 and again unfreeze all layers. We run for a total of 12 epochs, with initial learning rate  $10^{-3}$  and decay after 8 and 10 epochs. Due to this being the conventional BTDFactorisation strategy, we often refer to this as the *true* accuracy of a configuration.

## 5 Results

In this section, we demonstrate the effectiveness of GroSS. First, we explore group-size selection for network acceleration through search on our GroSS decomposition. Secondly, we justify the design of GroSS over a more simply using partial training schedules for group configurations.

### 5.1 Group-size Search

Here, we evaluate the performance of our search. We split our results by dataset, with our CIFAR-10 results being followed by our results on ImageNet.

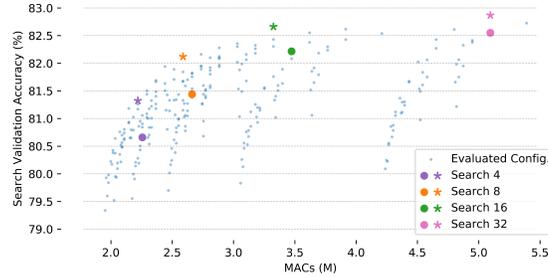


Fig. 3: Exhaustive search for the 4-layer network on CIFAR-10. Each search is colour coded. The circles and stars mark the performance of the baseline and found configuration, respectively.

For each search, we report the change in accuracy of the found configuration over the baseline, as well as the percentage reduction in MACs compared to the baseline. We are primarily concerned with exploring the impact of the number of groups on the performance of the network, rather than other design parameters such as the bottleneck dimensions. In our experimental setup, the inference cost between configurations varies only in the number of groups in each bottleneck. However, there is significant overhead from other layers in the network that remains constant between configurations. Therefore, we report the reduction in total MACs, as well as the contribution from the group layers alone. This provides greater insight into grouped architecture design and the performance of the searches using GroSS.

Table 2: Exhaustive search for our 4-layer network on CIFAR-10

4-Layer Network Configuration	MACs	Accuracy			$\Delta$ MACs		
		GroSS	True	$\Delta$ Acc.	Total	G.Conv	
Full	5.13M	-	83.99 (0.53)	-	-	-	
Baseline: $32 \ 32 \ 32$	<b>5.09M</b>	82.55 (0.07)	83.70 (0.05)	-	-	-	
$32 \ 16 \ 64$	<b>5.09M</b>	<b>82.87 (0.14)</b>	<b>84.05 (0.07)</b>	$\uparrow 0.35$	0.00%	0.00%	
Baseline: $16 \ 16 \ 16$	3.47M	82.22 (0.10)	82.94 (0.06)	-	-	-	
VBMF [18]: $16 \ 8 \ 16$	<b>3.33M</b>	81.76 (0.10)	82.83 (0.10)	$\downarrow 0.11$	$\downarrow 4.25\%$	$\downarrow 9.09\%$	
$8 \ 16 \ 64$	<b>3.33M</b>	<b>82.66 (0.11)</b>	<b>83.88 (0.10)</b>	$\uparrow 0.94$	$\downarrow 4.25\%$	$\downarrow 9.09\%$	
Baseline: $8 \ 8 \ 8$	2.66M	81.44 (0.16)	82.86 (0.10)	-	-	-	
$2 \ 16 \ 32$	<b>2.59M</b>	<b>82.12 (0.16)</b>	<b>83.50 (0.07)</b>	$\uparrow 0.64$	$\downarrow 2.77\%$	$\downarrow 9.09\%$	
Baseline: $4 \ 4 \ 4$	2.26M	80.66 (0.13)	82.37 (0.11)	-	-	-	
$1 \ 8 \ 16$	<b>2.22M</b>	<b>81.32 (0.16)</b>	<b>82.45 (0.15)</b>	$\uparrow 0.08$	$\downarrow 1.63\%$	$\downarrow 9.09\%$	
Depthwise: $1 \ 1 \ 1$	1.95M	79.34 (0.14)	81.70 (0.32)	-	-	-	

Table 3: Breadth-first search on our VGG-16 network on CIFAR-10

VGG-16 (CIFAR) Configuration	Accuracy			$\Delta$ MACs		
	MACs	GroSS	True	$\Delta$ Acc.	Total	G.Conv
Full	314M	-	91.52	-	-	-
<b>32 32 32 32 32 32 32 32 32 32 32 32 32</b>	<i>121M</i>	<i>90.97</i>	<b>91.57</b>	-	-	-
VBMF [18]	118M	90.97	91.31	$\downarrow 0.26$	$\downarrow 2.68\%$	$\downarrow 6.18\%$
<b>32 4 16 32 16 16 32 16 4 16 32 1</b>	<b>103M</b>	<b>91.31</b>	91.41	$\downarrow 0.16$	$\downarrow 14.6\%$	$\downarrow 33.7\%$
<b>16 16 16 16 16 16 16 16 16 16 16 16 16</b>	<i>94.6M</i>	<i>91.13</i>	<i>91.19</i>	-	-	-
<b>4 4 16 32 16 16 1 32 4 32 16 1</b>	<b>86.7M</b>	<b>91.28</b>	<b>91.31</b>	$\uparrow 0.12$	$\downarrow 8.36\%$	$\downarrow 30.2\%$
<b>4 4 4 4 4 4 4 4 4 4 4 4</b>	<i>74.9M</i>	<i>90.43</i>	<i>90.90</i>	-	-	-
<b>1 1 1 16 16 4 1 1 4 4 1 4</b>	<b>74.1M</b>	<b>90.97</b>	<b>91.14</b>	$\uparrow 0.24$	$\downarrow 1.11\%$	$\downarrow 12.6\%$
<b>1 1 1 1 1 1 1 1 1 1 1 1</b>	<i>70.0M</i>	<i>90.24</i>	<i>90.66</i>	-	-	-

**CIFAR-10.** The results of the exhaustive search on the 4-layer network are shown in Table 2, where the decomposition and tune is performed 5 times for each configuration and the mean and standard deviation are reported decompositions with uniform rank values across layers (4, 8, 16, and 32) are chosen as the baseline configurations for the search, such that we perform search across the range of possible configurations. For each baseline configuration we are able to find an alternative that is more accurate whilst requiring fewer operations. The results of the search are also visualised in Figure 3.

In the case of our CIFAR-10 VGG-16 network, the  $4^{12}$  configurations produced by our GroSS decomposition are too many to feasibly enable exhaustive evaluation. We, therefore, perform a breadth-first search. The full details of how this search is performed are described in Section 3.4. Results for this search on VGG-16 are shown in Table 3.

For the searches on the 4, 8 and 16 baselines, we are able to find configurations which meet the objective. However, in the case of the search below the 32 baseline, the found configuration’s true accuracy is less than that of the baseline. We speculate that this is because 32 is the maximum rank in the decomposition. Therefore, the rank of each layer can never be increased above that of the baseline. This means that configurations have less room to manoeuvre in targeting more heavy-duty layers at key stages of the network.

We also include a configuration found through Variational Bayesian Matrix Factorisation (VBMF) [18], which is used for one-shot rank selection in [11]. For both networks, we were able to find more accurate configurations which require fewer or the same number of operations than the VBMF rank selection. In fact, Kim *et al.* [11] note that, although they achieve good network compression results with the result of VBMF, they had not investigated whether this method of rank selection was optimal. The results in Table 2 demonstrate that VBMF is not optimal in this case, and GroSS is an effective tool to determine this.

Table 4: Breadth-first search for VGG-16 and ResNet-18 on ImageNet. \* denotes the configuration is using the decomposition structure from [26]

ImageNet Configuration	Accuracy			$\Delta$ MACs	
	MACs	GroSS	True $\Delta$ Acc.	Total	G.Conv
<b>VGG-16 (Full)</b>	15.49B	-	71.59	-	-
<b>16 16 16 16 16 16 16 16 16 16 16 16</b>	4.75B	70.25	70.77	-	-
<b>1 32 32 32 16 1 32 16 32 4 32 16</b>	<b>4.70B</b>	<b>70.40</b>	<b>70.82</b>	$\uparrow$ 0.04	$\downarrow$ 3.65%
<b>4 4 4 4 4 4 4 4 4 4 4 4</b>	3.78B	69.73	70.51	-	-
<b>1 1 4 4 4 4 32 16 1 1 32 1</b>	<b>3.78B</b>	<b>69.97</b>	<b>70.63</b>	$\uparrow$ 0.12	$\downarrow$ 1.69%
<b>1 1 1 1 1 1 1 1 1 1 1 1</b>	3.54B	68.98	70.28	-	-
[26]: <b>11 10 14 9 15 16 16 29 33 56 56 56*</b>	1.16B	62.64	66.85	-	-
<b>11 10 14 9 15 32 64 58 3 56 7 7*</b>	<b>1.16B</b>	<b>63.11</b>	<b>67.22</b>	$\uparrow$ 0.37	$\downarrow$ 4.89%
<b>ResNet-18 (Full)</b>	1.82B	-	69.76	-	-
<i>Baseline 16s</i>	738M	60.77	65.80	-	-
<b>16 16 4 32 32 32 32 32 16 32 32 32 16 4 32 16</b>	<b>715M</b>	<b>61.25</b>	<b>65.84</b>	$\uparrow$ 0.04	$\downarrow$ 3.18%
<i>Baseline 4s</i>	586M	60.02	65.46	-	-
<b>1 4 4 1 16 4 1 4 32 4 4 4 16 4 4 4</b>	<b>585M</b>	<b>60.31</b>	<b>65.44</b>	$\uparrow$ 0.18	$\downarrow$ 0.88%
<b>1 1 1 1 1 1 1 1 1 1 1 1 1</b>	547M	58.61	65.16	-	-

**ImageNet.** We now move to a larger, more complex dataset in ImageNet. We perform the same GroSS decomposition and breadth-first search on conventional VGG-16 and ResNet-18 structures. We search against baseline configurations of uniform 4s and 16s. The results are listed in Table 4 and we provide visualisation of the search on VGG-16 in Figure 4.

We also include results for an alternative decomposition structure and group configuration identical to that used in [26], which we detail in the supplementary material. This decomposition structure aggressively reduces widths in bottleneck layers to achieve a large compression ratio. In our search, we are able to show that the original configuration of groups within this structure is not optimal,

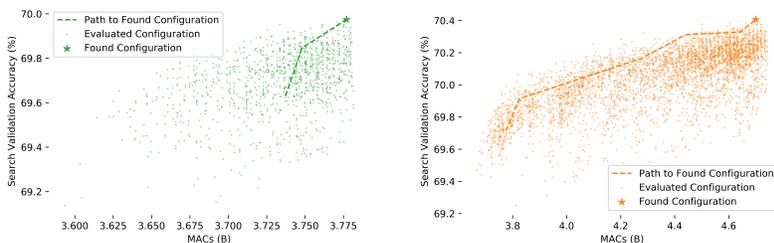


Fig. 4: Visualisation of the VGG-16 breadth-first search on ImageNet. (Left) and (Right) are the searches against the 4s and 16s baselines, respectively.

with our found configuration leading to a significant improvement in accuracy as well as a slight speed up.

The results show that exploration of group selection with GroSS generalises well across datasets and architectures. In every search performed, we found configurations that met the objective of increased accuracy with lower inference cost.

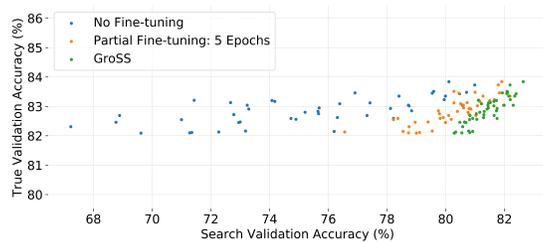


Fig. 5: Search-space vs true validation accuracy for our 4-layer network on CIFAR-10. Here we plot the accuracy of 45 random configurations of our 4-layer network for 3 different methods of obtaining a search-space. The accuracy of each configuration is plotted against its true validation accuracy.

## 5.2 GroSS Vs. Conventional Fine-tuning

In this section, we justify the need for GroSS by evaluating it against using a partial fine-tuning strategy for each individual configuration. For this, we select 45 random group-size configurations of our 4-layer network and fine-tune according to our individual schedule, which is outlined in Sec. 4.4, giving us a true validation accuracy for each configuration. We can then evaluate the validation accuracy of these same configurations in our GroSS search-space. This procedure allows us to visualise how representative GroSS is of the true accuracy.

For comparison, we also include the validation accuracy of the same configurations with no fine-tuning, as well as with a shortened schedule of 5 epochs. The partial fine-tune could be considered a reasonable solution to reducing the burden of training while performing a configuration search.

We visualise the search-space against the true validation accuracies in Figure 5. Qualitatively, it can be seen that the validation accuracies produced by GroSS produce a significantly more consistent search space. The points appear to be more tightly distributed and closer to the ideal distribution ( $y = x$ ). To measure this quantitatively, we compute the top-5 average precision of the search spaces. We simulate searches across the entire range of configurations by evaluating the average precision at multiple slices through the search-space. This allows for comparison across the space, not just the most accurate group configurations.

Table 5: Average precision across the range of the search space. We compute the average precision using the top 5 true validation accuracies as positive recalls. “X ↓” refers to the average precision computed after the top-X configurations have been removed from the search

Fine-tune Strategy	Average Precision (Top 5)				
	All	10 ↓	20 ↓	30 ↓	Mean
No Fine-tuning	43.0	43.8	86.3	69.8	60.7
Partial Fine-tuning	35.2	<b>64.0</b>	42.5	66.4	52.0
GroSS	<b>44.1</b>	63.8	<b>94.3</b>	<b>82.5</b>	<b>71.2</b>

Table 5 lists the results of this average precision computation. GroSS is consistently as good or better than no fine-tuning and the partial schedule at each slice. This leads to a significant improvement in search performance across the range of configurations which is highlighted by the mean average precision.

When making the comparison between GroSS and a partial training strategy, it is worth considering the computational requirements of each. Running inference for a new configuration in either of the conventional decompositions requires a new network to be initialised, and weights to be loaded. However, since group-sizes are handled dynamically within a GroSS decomposition, switching between them is essentially free, with no structure change or weight loading. This leads to GroSS having a significant speed improvement for running inference (7s vs 287s) over the 45 configurations. This only increases with more configurations tested. For example, the inference for the exhaustive search on the GroSS decomposition of the 4-layer network took only 9s for 252 configurations. Similarly, the total number of training epochs for the partial training strategy increase linearly with the number of configurations, but remain constant for GroSS. With larger search-spaces, such as those visualised in Figure 4, the accuracy and performance benefits of GroSS combine to make grouped architecture search feasible where it might not have been before.

## 6 Conclusions

In this paper, we have presented GroSS, a series BTDFactorisation which allows for the dynamic assignment and simultaneous training of differing numbers of groups within a layer. We have shown how GroSS-decomposed layers can be combined to train an entire grouped convolution search space at once. We confirmed the value of these configurations through an exhaustive search and a breadth-first search. We further demonstrate that, without GroSS, these searches would be less effective and dramatically less efficient.

**Acknowledgements** We gratefully acknowledge the European Commission Project Multiple-actOrs Virtual Empathic CAREgiver for the Elder (MoveCare) for financially supporting the authors for this work.

## References

1. Brock, A., Lim, T., Ritchie, J.M., Weston, N.: SMASH: One-shot model architecture search through hypernetworks. arXiv Preprint arXiv:1708.05344 (2017)
2. Chen, Y., Jin, X., Kang, B., Feng, J., Yan, S.: Sharing residual units through collective tensor factorization to improve deep neural networks. In: IJCAI. pp. 635–641 (2018)
3. 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)
4. De Lathauwer, L.: Decompositions of a higher-order tensor in block terms—part II: Definitions and uniqueness. *SIAM Journal on Matrix Analysis and Applications* **30**(3), 1033–1066 (2008)
5. 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)
6. Elthakeb, A.T., Pilligundla, P., Yazdanbakhsh, A., Kinzer, S., Esmailzadeh, H.: ReLeQ: A reinforcement learning approach for deep quantization of neural networks. arXiv Preprint arXiv:1811.01704 (2018)
7. 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)
8. 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)
9. Ioannou, Y., Robertson, D., Cipolla, R., Criminisi, A.: Deep roots: Improving cnn efficiency with hierarchical filter groups. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 1231–1240 (2017)
10. Jaderberg, M., Vedaldi, A., Zisserman, A.: Speeding up convolutional neural networks with low rank expansions. arXiv Preprint arXiv:1405.3866 (2014)
11. Kim, Y.D., Park, E., Yoo, S., Choi, T., Yang, L., Shin, D.: Compression of deep convolutional neural networks for fast and low power mobile applications. arXiv Preprint arXiv:1511.06530 (2015)
12. Kossaifi, J., Panagakis, Y., Anandkumar, A., Pantic, M.: Tensorly: Tensor learning in python. *The Journal of Machine Learning Research* **20**(1), 925–930 (2019)
13. Krizhevsky, A., Nair, V., Hinton, G.: The CIFAR-10 dataset. Online: [Http://www.Cs.Toronto.Edu/Kriz/Cifar](http://www.Cs.Toronto.Edu/Kriz/Cifar). HTML **55** (2014)
14. 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)
15. Lebedev, V., Ganin, Y., Rakhuba, M., Oseledets, I., Lempitsky, V.: Speeding-up convolutional neural networks using fine-tuned CP-decomposition. arXiv Preprint arXiv:1412.6553 (2014)
16. Li, L., Talwalkar, A.: Random search and reproducibility for neural architecture search. arXiv Preprint arXiv:1902.07638 (2019)
17. Liu, H., Simonyan, K., Yang, Y.: DARTS: Differentiable architecture search. arXiv Preprint arXiv:1806.09055 (2018)
18. Nakajima, S., Sugiyama, M., Babacan, S.D., Tomioka, R.: Global analytic solution of fully-observed variational Bayesian matrix factorization. *Journal of Machine Learning Research* **14**(Jan), 1–37 (2013)

19. 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)
20. Sifre, L., Mallat, S.: Rigid-motion scattering for image classification (2014)
21. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. arXiv Preprint arXiv:1409.1556 (2014)
22. Tan, M., Chen, B., Pang, R., Vasudevan, V., Sandler, M., Howard, A., Le, Q.V.: MnasNet: Platform-aware neural architecture search for mobile. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 2820–2828 (2019)
23. Tan, M., Le, Q.V.: EfficientNet: Rethinking model scaling for convolutional neural networks. arXiv Preprint arXiv:1905.11946 (2019)
24. Tucker, L.R.: Some mathematical notes on three-mode factor analysis. *Psychometrika* **31**(3), 279–311 (1966)
25. Vanhoucke, V., Senior, A., Mao, M.Z.: Improving the speed of neural networks on CPUs (2011)
26. Wang, P., Hu, Q., Fang, Z., Zhao, C., Cheng, J.: DeepSearch: A fast image search framework for mobile devices. *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)* **14**(1), 6 (2018)
27. Wu, B., Dai, X., Zhang, P., Wang, Y., Sun, F., Wu, Y., Tian, Y., Vajda, P., Jia, Y., Keutzer, K.: FBNet: Hardware-aware efficient convnet design via differentiable neural architecture search. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 10734–10742 (2019)
28. 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)
29. 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)
30. 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)