

# Filter redistribution templates for iteration-less convolutional model reduction

Anonymous authors

Paper under double-blind review

## ABSTRACT

Automatic neural network discovery methods face an enormous challenge caused for the size of the search space. A common practice is to split this space at different levels and to explore only a part of it. Neural architecture search methods look for how to combine a subset of layers, which are the most promising, to create an architecture while keeping a predefined number of filters in each layer. On the other hand, pruning techniques take a well known architecture and look for the appropriate number of filters per layer. In both cases the exploration is made iteratively, training models several times during the search. Inspired by the advantages of the two previous approaches, we proposed a fast option to find models with improved characteristics. We apply a small set of templates, which are considered promising, for make a redistribution of the number of filters in an already existing neural network. When compared to the initial base models, we found that the resulting architectures, trained from scratch, surpass the original accuracy even after been reduced to fit the same amount of resources.

## 1 INTRODUCTION

Convolutional neural networks are built by stacking layers of neurons following the principle explained by Fukushima’s Neocognitron model (Fukushima, 1980). The Neocognitron design made neural networks invariant to shift in feature locations by arranging cells locally connected in a hierarchical architecture.

Instead of connecting every neuron of the previous layer to all the neurons in the next layer, convolutional networks connections are only made to a small region. Given that different regions share the same weights, the layer can be implemented as a convolution operation using the set of shared weights known as kernel. The complete input is processed by shifting the kernel at uniform steps, normally overlapping parts of the input. To improve shift invariance, convolutional networks needs to rely less in the exact position of a feature and a simple solution is to have a lower resolution performed by averaging the values of neighbouring points in the image in a operation known as spatial subsampling.

An important consideration to create a convolutional network model is the number of filters, required at every layer. The Neocognitron implementation for example, keeps a fixed number of filters for each layer in the model. A very common practice has been to use a by-pyramidal architecture. The number of filters across the different layers is usually increased

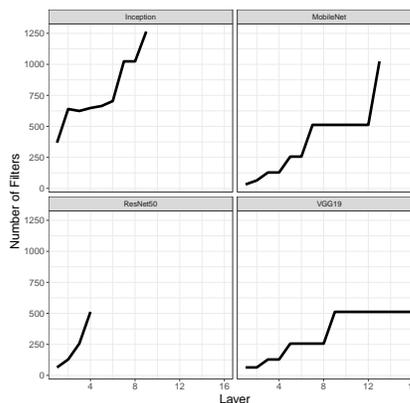


Figure 1: Increasing filters per layer is a common design in convolutional models. We apply templates to existing models to change filters distribution. (ResNet shows filters per sets of residual blocks. The number is keep fixed inside them).

as the size of the feature maps decrease. This pattern was first proposed by LeCun et al. (1998) with the introduction of LeNet and can be observed in a diverse set of models such as VGG, ResNet and MobileNet (Figure 1). Even models obtained from automatic model discovery, like NASNet, follow this principle inasmuch as neural network methods are mainly formulated to search for layers and connections.

It can be found in (LeCun et al., 1998) that the reason behind this progressive increase in the number of kernels is to compensate a possible loss of the representation caused by the spatial resolution reduction. In recent models, what seems to be the real reason is a practical issue (Chu & Krzyżak, 2014), to improve performance by keeping a constant number of operations in each layer.

The Pyramidal distribution of filters has perpetuated over two areas of model discovery. The methods in Automatic Neural Architecture Search (Liu et al., 2018a; Tan et al., 2019) explore the models built by combining a predefined set of layers, commonly with a pyramidal distribution of filters. On the other side, Network Pruning aims to reduce a model computational resources demands by selecting and removing weights that match some rule, commonly the closest to zero values, but starting from models that present this pyramidal distribution.

To the best of our knowledge, it remains unknown if this pyramidal distribution of filters is also beneficial to different aspects of model performances other than the number of operations (such as accuracy, memory footprint, inference time and model compression level).

This paper explores on the topic by comparing models against versions of themselves with the same general structures but with the distribution of filters across the layers changed. We present a fast to implement method for model discovery, that take the goal of structured pruning methods, that is finding the lowest number of filters for a model while maintaining accuracy. Our exploration technique only tests over a small subset of diverse filter distributions, called templates, and then reduces the model proportionally to match some resource budget. Our experiments show that by using our proposed templates, the resulting models keep comparable accuracy as original models in classification tasks but they present reductions in the number of parameters and memory footprint.

The contributions of this paper are the following: 1) it provides evidence that pyramidal distribution of filters in convolutional network models is usually optimised for a distributed GPU operation across layers, and simple changes to that distribution leads to improvements in metrics such as number of parameters or memory footprint; 2) it highlights that most recent models, which have had a more detailed tuning in the filter distribution, presents resiliency in accuracy to changes in the filter distribution, a phenomena that requires further research and explanation; 3) it shows that redistributing filters in a model and then applying a width multiplier operation can be seen as a pruning technique which produces smaller models than just applying the width multiplier to the original models; 4) it gives to classical models a repositioning of their value when measuring others resources equally important for practical implementation.

The rest of the paper is structured as follows: chapter 2 explores the most recent methods to reduce the size of neural network architectures. Chapter 3 describes the set of templates for filter distribution and how to implement this change in a convolutional network model. Chapter 4 compares dissimilar allocations of filters and their effect in model performances. Finally, chapter 5 briefly explains the findings derived from experiments.

## 2 RELATED WORK

The process of designing a Neural Networks is a task mainly based on experience and experimentation that consumes a lot of time and computational resources. With the increase in the use of Neural Networks, and particularly Convolutional Networks for computer vision problems, a mechanism to automatically find the best architecture has become a requirement in the field of Deep Learning. Although some works have been published several years ago (Pinto et al., 2009; Kuri-Morales, 2014) trying to solve the topic of automatic architecture generation, they had not provided competitive results compared to hand-crafted architectures. However, current works start to lead the state of the art models (Elsken et al., 2018).

But even in this automatic methods, one key feature that constantly has been adopted is the selection of the number of filters in each layer in the final model. The filters are set in such a way to have an increasing number as the layers go deeper. Pruning methods have done some work in this field but with the belief that the weights obtained at the end of the training process are important to the pruning method.

One common characteristic of many model discovery methods is that the search process is very time consuming, normally in the order of thousands of GPU days. In this sense, a remarkable improvement is presented by Liu et al. (2018b). The method uses a relaxation condition to transform the selection of layers in the architecture to a continuous space using a *softmax* function. The relaxation allows to perform simultaneous search of weights and architecture using gradient descent. The method converges after one day of GPU time.

On the side of Pruning methods the search also involves training models for several iterations to select the correct weights to remove (Frankle & Carbin, 2018; He et al., 2019), or at least increasing the computation during the training when doing jointly training and search (Leclerc et al., 2018). Recently, Liu et al. (2018c) suggest that accuracy obtained by pruning techniques can be reached by training from scratch.

Our work relates to (Gordon et al., 2018) in the sense that their method it is not restricted to reducing filters but also to increase them to see if the increment is beneficial. Ours differs because it does not requires to train the model other than in the final stage, after making some predefined changes to the number of filters using templates.

### 3 FILTER DISTRIBUTION TEMPLATES

Recent pruning methods have shown different filter distribution patterns emerging when reducing VGG models that defied the pyramidal design as the best distribution for a model. Instead of finding the adequate number of filters with automatic pruning techniques, we propose to adjust the filters of a convolutional network model with a small number of templates depicted in figure 2 that could be beneficial for model performances other than the number of operations (such as accuracy, memory footprint and inference time and model compression level). In particular, we test a fixed number of filters as the original Neocognitron design, but also, inspired by the patterns found by the actual pruning methods, we test also a quadratic distribution.

We define a convolutional neural network base model as a set of layers  $L = 1, \dots, D + 1$  each of them containing a number of filters  $f_l$  and a total number of filters  $F = \sum_{l=1}^D f_l$ . We want to test if the common heuristic of distributing  $F$  having  $f_{l+1} = 2f_l$  each time the feature map is halved, is advantageous to the model over other distributions of  $f_l$  when evaluating performance, memory footprint and inference time.

It should be noticed that that final layer  $D + 1$  remains with the same number of filters according to the classification task under evaluation, therefore it is not taken into account in the equations. Another important consideration is that, in architectures composed for modules or blocks (ResNet and Inception), it is easier to change the number of filters in the module as a whole than to change the filters in each particular layer inside the module and then to ensure that concatenations and additions from the previous layers match the correct number of filters. We have adopted then, this assumption of taking blocks as single layers to set the values of  $f_l$ .

**Uniform Template.** The most immediate distribution to evaluate is, as the original Neocognitron, an uniform distribution of filters. Computing the number of filters in an uniform distribution is straightforward. Adding up the filters in each layer from the base model and divide them by the number of layers give the number of filters to be set in each of them. Formally, we compute the new number in each layer as  $f'_l = F/D \quad \forall l \in \{1, \dots, D\}$ .

In this way, changing the distribution for a VGG19 model built exclusively with sixteen convolutional layers, one final unchangeable fully connected layer and a total number of filters of  $F = 5504$  produces a model with  $f'_l = 5504/16 = 344$  filters in each layer.

**Reverse Template.** Another naive transformation for the filter distribution adopted in this paper is reversing the filters in every layer. Our final model with this template is defined by the filters  $f'_l = f_{D-l+1}$ .

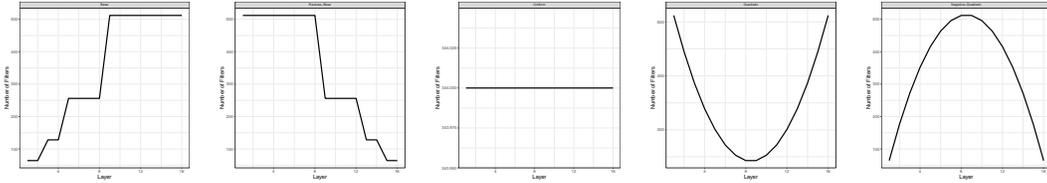


Figure 2: Filters per layer using the proposed templates for filter redistribution in the VGG 19 model. Base distribution, which is the original distribution, shows the common design of growing the filters when resolution of feature maps decreases in deeper layers.

**Quadratic Template.** The third distribution we evaluated is characterised by a quadratic equation  $f'_l = al^2 + bl + c$  and consequently, has a parabolic shape with the vertex in the middle layer. We set this layer to the minimal number of filters in the base model  $f'_{D/2} = \min(f_l) \quad l \in \{1, \dots, D\}$  so, the number of filters is described by  $f'_{D/2} = f_{min}$ . Also, we find the maximum value in both the initial and final convolutional layers, thus  $f'_1 = f'_D$ .

To compute the new number of filters in each layer we solve the system of linear equations given by *i*) the restriction of the total number of filters in  $\sum_{l=1}^D (f'_l) = \sum_{l=1}^D (al^2 + bl + c) = F$ , that can be reduced to  $\left(\frac{D^3}{3} + \frac{D^2}{2} + \frac{D}{6}\right)a + \left(\frac{D^2}{2} + \frac{D}{2}\right)b + Dc = F$ , *ii*) the equation produced by the value in the vertex  $f'_{D/2} = \frac{D^2}{2}a + \frac{D}{2}b + c = f_{min}$  and *iii*) the equality from the maximum values which reduces to  $(D^2 - 1)a + (D - 1)b = 0$ .

**Negative Quadratic Template.** The final template is also a parabola but with the vertex in a maximum, that is, a negative quadratic curve. The equation is still the same quadratic equation that the previous template but the restrictions change. Instead of defining a value in the vertex,  $f'_l$  at the initial and final convolutional layers are set to the minimal number of filters in the base model  $f'_l = f_{min} \quad l \in \{1, D\}$ . The number of filters in each layer is computed again with a system of equations specified by *i*) the restriction of the total number of filters as in the quadratic template, and the two points already known in the first and last convolutional layers defined by *ii*)  $a + b + c = f_{min}$  and *iii*)  $D^2a + Db + c = f_{min}$ .

Once the model has been readjusted with the new number of filters per layer we use a *width multiplier* to test different levels of model compression to make comparable evaluations given that the change in the distributions of kernels modifies the number of parameters, memory consumption and speed. The width multiplier only reduces or increases the set of new filters  $f'_l$  proportionally in every layer.

## DATASETS

We trained over two datasets traditionally used for convolutional network evaluation: CIFAR-10 and CIFAR-100 (Krizhevsky et al., 2009). Both datasets contain a train set of 50,000 images and a test set of 10,000 images with a resolution of 32x32 and three colour channels. They were published for classification tasks for ten and one hundred classes respectively.

## CONVOLUTIONAL NETWORK MODELS

The state-of-the-art networks evaluated represent some of the highest performing CNNs on the ImageNet challenge in the previous years (Russakovsky et al., 2015). They have been primarily tested on classification tasks and also have demonstrated a strong ability to generalise to images outside the ImageNet dataset. Therefore, it is expected they perform well in the CIFAR datasets.

The VGG network architecture (Simonyan & Zisserman, 2014) is recognised by its simplicity. It is composed by sequential convolutional layers followed by max-pooling reduction layers. The final classification is managed by fully-connected layer and a Softmax classifier. The main disadvantage

of these networks is the size of their parameters. In these paper we use the version of the model with just one fully connected layer in the final classification section.

ResNet (He et al., 2016) succeeds on the problem of training very deep CNNs by reformulating the assumption that the network blocks are modelling a function closer to an identity mapping than to a zero mapping. Therefore, it should be easier to find differences with reference to an identity rather than to a zero mapping. This assumption is carry out by adding additional references at the end of building blocks.

The Inception/GoogleNet architecture (Szegedy et al., 2015) make use of the Inception module conceived as a multi-level feature extractor allowing simultaneous extraction features of several sizes within the same module of the network.

The MobileNet network (Howard et al., 2017) is built on depthwise separable convolutions except for the first layer which is a full convolution. All layers are followed by a batch normalisation and ReLU nonlinearity with the exception of the final fully connected layer which consist of a softmax layer for classification.

#### 4 MODELS COMPARISON UNDER SIZE, MEMORY FOOTPRINT AND SPEED

In this section we first investigate the effects of applying different templates to the global distribution of kernels in well known convolutional neural network models (VGG, ResNet, Inception and MobileNet). We compare models under the basis of size, memory and speed in two of the popular datasets for classification tasks.

All experiments have models fed with images with the common augmentation techniques of padding, random cropping and horizontal flipping. Our experiments were run in a NVidia Titan X Pascal 12GB GPU adjusting the batch size to 64 to fit in GPU Memory at the maximum model scaling.

##### TEMPLATE EFFECT OVER THE BASELINE MODELS

We conducted a first experiment to test our proposed templates on the selected architectures. All convolutional models, with and without templates, were trained for 160 epochs using the same conditions: stochastic gradient descent (SGD) with a scheduled learning rate starting in 0.1 for the first 80 epochs, 0.01 for the next 40 epochs and finally 0.001 for the remaining epochs.

	Model	Base	Redistribution Templates			
			Reverse-Base	Uniform	Quadratic	Negative-Quadratic
CIFAR-10	VGG-19	93.7	94.59	94.57	<b>94.64</b>	94.15
	ResNet-50	94.6	95.37	95.32	94.44	<b>95.42</b>
	Inception	<b>94.92</b>	94.63	94.82	94.82	94.63
	MobileNet	91.35	<b>93.26</b>	93.2	90.92	92.33
CIFAR-100	VGG-19	72.38	<b>75.05</b>	74.44	74.05	74.46
	ResNet-50	76.96	75.01	<b>77.44</b>	75.37	76.84
	Inception	<b>78.76</b>	77.61	78.52	78.49	76.73
	MobileNet	67.98	70.54	<b>72.04</b>	66.41	71.18

Table 1: Model performances with the original distribution and four templates for the same number of filters evaluated on CIFAR-10 and CIFAR-100 datasets. After filter redistribution, most models surpass the base accuracy.

The results are presented in Table 1. It is shown that for VGG, ResNet and MobileNet, the model accuracy improves when the template is applied excepting the Inception architecture, which presents the highest accuracy of all base models in both datasets. Even though, templates are only able to change Inception accuracy in less than 1%. This is surprising given the drastic modifications that the model is suffering after the change of filter distribution. Models that share a sequential classical architecture such as VGG and MobileNet, show a better improvement when using templates.

	Model	Base	Reverse-Base	Uniform	Quadratic	Negative-Quadratic
Parameters (Millions)	VGG-19	20.0	20.0	16.0	<b>15.8</b>	20.0
	ResNet-50	23.5	23.1	<b>12.9</b>	19.0	33.0
	Inception	6.2	6.7	6.9	7.2	<b>3.2</b>
	MobileNet	3.2	<b>2.2</b>	2.9	3.2	3.6
Memory Footprint (GB/batch)	VGG-19	<b>1.7</b>	10.7	5.2	9.3	1.9
	ResNet-50	10.2	23.0	9.6	19.4	<b>8.7</b>
	Inception	5.7	10.8	7.0	8.0	<b>3.7</b>
	MobileNet	2.9	6.8	3.3	7.7	<b>2.0</b>
Inference Time ( $\mu$ s)	VGG-19	<b>73.8</b>	432.8	102.9	366.6	95.2
	ResNet-50	<b>645.4</b>	1,968.6	1,263.1	1,350.0	1,324.0
	Inception	707.8	1,422.7	1,093.9	1,187.5	<b>486.8</b>
	MobileNet	<b>74.5</b>	488.2	560.1	492.0	89.0

Table 2: Parameters, memory and inference time for selected models when applying our templates keeping the same number of filters evaluated on the CIFAR-10 dataset. Models are normally optimised to fast GPU operation, therefore the original base distribution has a good effect in speed but the redistribution of filters induced by our templates makes models capabilities improve on the other metrics.

When evaluated under other metrics (Table 2), models are affected differently with each template and model. The Reverse-Base, Uniform and Quadratic templates show some reductions in the number of parameters while Negative Quadratic template reduces the memory usage for all models except VGG. Inference Time is affected negatively for the templates. This is an expected result as original models are designed to perform well in the GPU. The Inception model, which accuracy only decreases with templates, shows a remarkable improvement in speed with a reduction of 30% over inference time respect to the base model. It is important to notice that a reduced number of parameters does not correspond to a low consumption of memory, not even a small inference time. Some of the causes are the difference in feature map resolution for filters in different layers, the need to keep early feature maps in memory for late layers and the restrictions for improving parallelisation in the computational graph of the model.

#### TEMPLATE EFFECT WITH SIMILAR RESOURCES

It can be argued that models obtained with templates make use of more resources such as memory or number of operations in the GPU (reflected in the low inference speed). So, we formulated a second experiment that makes proportional changes in the models after applying the templates. We not only apply reductions to the models but also increments in order to observe if the actual total number of filters is adequate for the task the model is performing or if the model accuracy could improve by adding more filters. Thus, we create curves for each template applying an uniform reduction using a width multiplier with values of 1.6, 1.3, 1.0, 0.8, 0.5, 0.25, 0.1 and 0.05. These curves of reduction allow comparison under the same amount of resources as well as compares the use of resources under the same accuracy. The experiment also shows the level of reduction that our models can tolerate without a significant loss in accuracy.

We add dashed lines to every plot to be used as a reference for the model with the original distribution and no reductions which is the point where both vertical and horizontal dashed lines cross. In general, any arbitrary vertical line in the plot compares accuracy between models with amount of resources (parameters, memory or speed). On the other side, any arbitrary horizontal line compares the resources taken for each model under each template to produce similar accuracy.

Evaluating a model performance using accuracy and parameters is by far the default approach. We show models performances with these metrics in figure 3. VGG and MobileNet models improve in accuracy almost with any template in CIFAR-10 and CIFAR-100. Under reductions, their original accuracy can be reached with less than 25% of the original parameters in the two models. ResNet shows more small improvement when compare similar resource consumption but it can be seen that templates make the model able to being reduced more before a drop in accuracy. Inception behavior

considering the same resources remains similar no matter the template used. In general, the uniform template seems to get the best parameter efficiency for all the models.

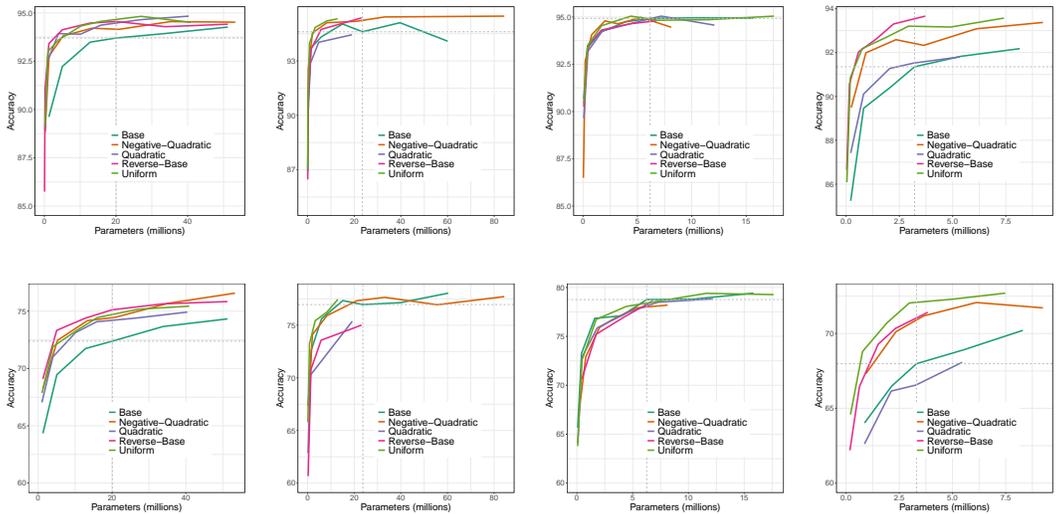


Figure 3: Accuracy versus Parameters in CIFAR-10 (top row) and CIFAR-100 (bottom row) datasets using templates with VGG, ResNet, Inception and MobileNet. Curves are created by reducing models using width multiplier scaling. An arbitrary vertical line in the plot compares templates effects using the same amount of parameters. An arbitrary horizontal line compares parameters for reaching the same accuracy.

Figure 4 presents an evaluation considering memory footprint and accuracy. We are convinced that for practical implementations, comparing parameters is not a good option. Table 2 has shown that models with a small number of parameters are not necessary related with a small memory footprint or bigger speed. We observe again that VGG and MobileNet accuracy is enhanced by templates. More than 50% of memory consumption can be reduced in both models while producing the same accuracy. With this metric, ResNet and Inception improve slightly in CIFAR-10 with the Negative-Quadratic template but they perform lower with the rest of templates. We attribute the lower efficiency in memory to the fact that in all the templates but Negative-Quadratic, we are increasing the number of higher resolution feature maps, those more memory costly, that are kept at the initial layers.

One final comparison also important for practical issues is inference time. We present our results in figure 5 where it is observed the patten of improvement for VGG and MobileNet and a degradation of inference time when adopting templates in ResNet and Inception. In particular Inception shows an improvement with the Negative-Quadratic template in CIFAR-10.

By looking results in inference time it can look unpromising to apply templates. However we can take a different perspective, by sacrificing inference speed it is possible to obtain models with a better accuracy. This could be an unwanted decision but it is frequently taken. It is clearly stated in the inference time between different original models. For example by using ResNet the accuracy has improve compared to the obtained by VGG in the two datasets tested, but at the cost of increasing the time for inference. On the contrary, looking for speed enhancement MobilNet has sacrificed accuracy. In these sense, our templates are still competitive when compared to searching for a totally different model in order to improve accuracy.

## 5 CONCLUSIONS

The most common design of convolutional neural networks when choosing the distribution of the number of filters is to start with a few and then to increase the number in deeper layers. We challenged this design by evaluating some architectures with a variate set of distributions on the CIFAR

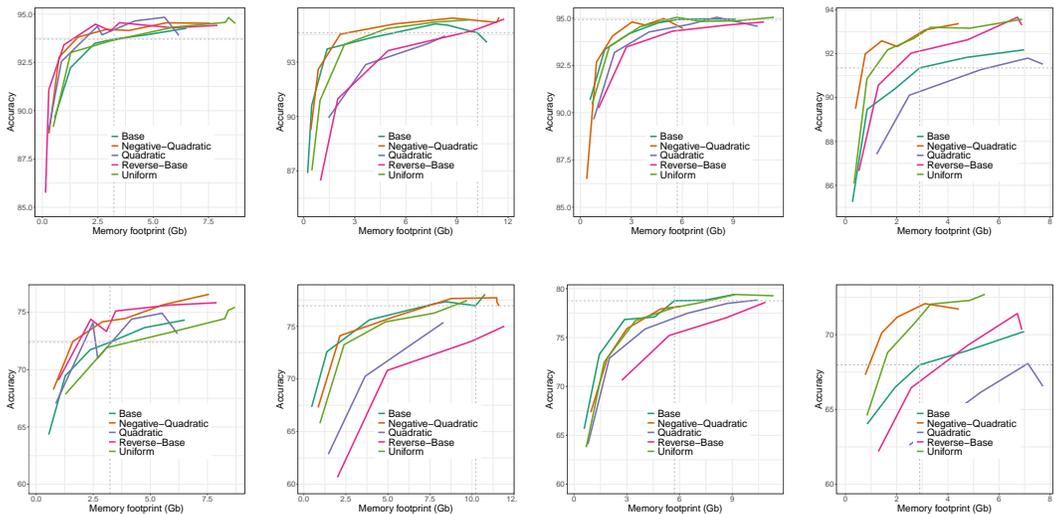


Figure 4: Accuracy versus Memory Footprint (batch size = 64) in CIFAR-10 (top row) and CIFAR-100 (bottom row) datasets using templates with VGG, ResNet, Inception and MobileNet. Curves are created by reducing models using width multiplier scaling. An arbitrary vertical line in the plot compares templates effects in accuracy using the same amount of memory. An arbitrary horizontal line compares memory consumption for reaching the same accuracy. Dashed lines crosses in the original model with no reduction.

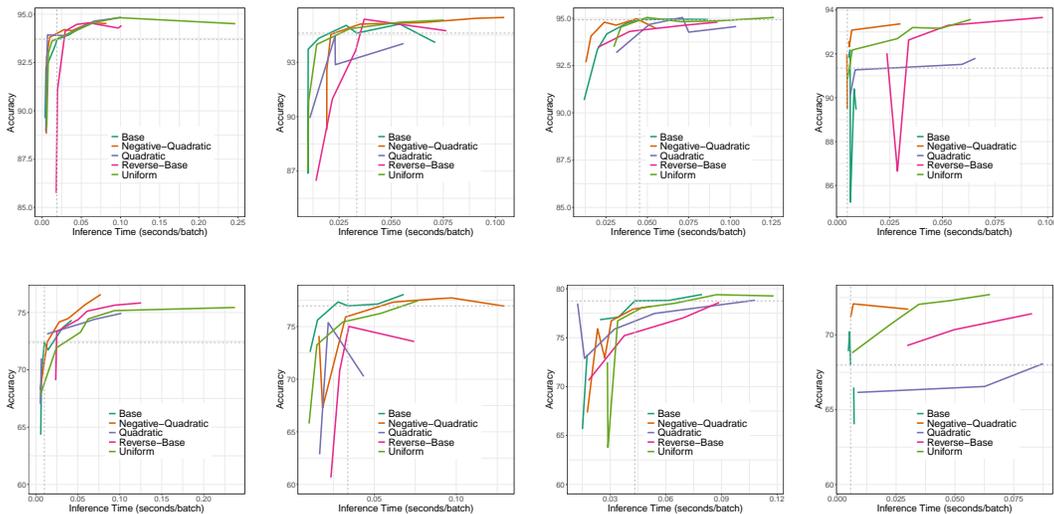


Figure 5: Accuracy versus Inference Time (batch size = 64) in CIFAR-10 (top row) and CIFAR-100 (bottom row) datasets using templates with VGG, ResNet, Inception and MobileNet. Curves are created by reducing models using width multiplier scaling. An arbitrary vertical line in the plot compares templates effects in accuracy between models with same inference speed. An arbitrary horizontal line compares inference time of models with the same accuracy.

datasets. Our results suggest that this pyramidal distribution of filters could not be a good option for obtaining the highest accuracy or even the highest parameter efficiency.

The method presented applies a set of templates for redistributing the number of filters originally assigned to each layer in existent convolutional network models and then trains the models from scratch. This redesign and the following proportional reduction can be achieved without any previous training process to select particular weights.

Our experiments show that the models, with the same amount of filters but a different distribution produced by our templates, improves significantly in accuracy. After being pruned uniformly, they can obtain the same accuracy than the original models using less resources such as number of parameters and memory footprint.

Results also reveals an interesting behaviour in the evaluated models: a strong resilience to changes in filter distribution. The variation in accuracy for all models after administering templates is less than 5% despite the modifications in the distributions are considerable. This finding strengthen our belief that it is not worth exploring the whole space of filter distributions to find the best solution at the cost of training models for a large number of iterations, as this solution possibly won't be too distant of the baseline performance. However, it is possible to explore just some distinct distributions such as those represented by our templates, that produces benefits depending on the resource to be optimised.

#### ACKNOWLEDGMENTS

Removed for blind review

#### REFERENCES

- Joseph Lin Chu and Adam Krzyżak. Analysis of feature maps selection in supervised learning using convolutional neural networks. In *Canadian Conference on Artificial Intelligence*, pp. 59–70. Springer, 2014.
- Thomas Elsken, Jan Hendrik Metzen, and Frank Hutter. Neural architecture search: A survey. *arXiv preprint arXiv:1808.05377*, 2018.
- Jonathan Frankle and Michael Carbin. The lottery ticket hypothesis: Finding sparse, trainable neural networks. *arXiv preprint arXiv:1803.03635*, 2018.
- Kunihiko Fukushima. Neocognitron: A self-organizing neural network model for a mechanism of pattern recognition unaffected by shift in position. *Biological cybernetics*, 36(4):193–202, 1980.
- Ariel Gordon, Elad Eban, Ofir Nachum, Bo Chen, Hao Wu, Tien-Ju Yang, and Edward Choi. Morphnet: Fast & simple resource-constrained structure learning of deep networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1586–1595, 2018.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- Yang He, Ping Liu, Ziwei Wang, Zhilan Hu, and Yi Yang. 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.
- 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.
- Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. Technical report, Citeseer, 2009.
- Angel Fernando Kuri-Morales. The best neural network architecture. In *Mexican International Conference on Artificial Intelligence*, pp. 72–84. Springer, 2014.

- Guillaume Leclerc, Manasi Vartak, Raul Castro Fernandez, Tim Kraska, and Samuel Madden. Smallify: Learning network size while training. *arXiv preprint arXiv:1806.03723*, 2018.
- Yann LeCun, Léon Bottou, Yoshua Bengio, Patrick Haffner, et al. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- 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 *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 19–34, 2018a.
- Hanxiao Liu, Karen Simonyan, and Yiming Yang. Darts: Differentiable architecture search. *arXiv preprint arXiv:1806.09055*, 2018b.
- Zhuang Liu, Mingjie Sun, Tinghui Zhou, Gao Huang, and Trevor Darrell. Rethinking the value of network pruning. *arXiv preprint arXiv:1810.05270*, 2018c.
- Nicolas Pinto, David Doukhan, James J DiCarlo, and David D Cox. A high-throughput screening approach to discovering good forms of biologically inspired visual representation. *PLoS computational biology*, 5(11):e1000579, 2009.
- Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision*, 115(3):211–252, 2015.
- Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1–9, 2015.
- Mingxing Tan, Bo Chen, Ruoming Pang, Vijay Vasudevan, Mark Sandler, Andrew Howard, and Quoc V Le. Mnasnet: Platform-aware neural architecture search for mobile. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2820–2828, 2019.