



HAL
open science

Learning Sparse Filters In Deep Convolutional Neural Networks With A l_1 / l_2 Pseudo-Norm

Anthony Berthelier, Yongzhe Yan, Thierry Chateau, Christophe Blanc, Stefan Duffner, Christophe Garcia

► **To cite this version:**

Anthony Berthelier, Yongzhe Yan, Thierry Chateau, Christophe Blanc, Stefan Duffner, et al.. Learning Sparse Filters In Deep Convolutional Neural Networks With A l_1 / l_2 Pseudo-Norm. CADL 2020: Workshop on Computational Aspects of Deep Learning - ICPR 2020, Jan 2021, Milan, Italy. 10.1007/978-3-030-68763-2_50 . hal-03048750

HAL Id: hal-03048750

<https://hal.science/hal-03048750>

Submitted on 9 Dec 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Learning Sparse Filters In Deep Convolutional Neural Networks With A l_1/l_2 Pseudo-Norm^{*}

Anthony Berthelie^{1*}[0000-0003-0545-0345], Yongzhe Yan¹[0000-0003-3294-7867],
Thierry Chateau¹[0000-0003-4854-5686], Christophe Blanc¹, Stefan
Duffner²[0000-0003-0374-3814], and Christophe Garcia²[0000-0001-7997-9837]

¹ Universite Clermont Auvergne, Institut Pascal, Clermont-Ferrand, France

² INSA Lyon, LIRIS, Lyon, France

Abstract. While deep neural networks (DNNs) have proven to be efficient for numerous tasks, they come at a high memory and computation cost, thus making them impractical on resource-limited devices. However, these networks are known to contain a large number of parameters. Recent research has shown that their structure can be more compact without compromising their performance.

In this paper, we present a sparsity-inducing regularization term based on the ratio l_1/l_2 pseudo-norm defined on the filter coefficients. By defining this pseudo-norm appropriately for the different filter kernels, and removing irrelevant filters, the number of kernels in each layer can be drastically reduced leading to very compact Deep Convolutional Neural Networks (DCNN) structures. Unlike numerous existing methods, our approach does not require an iterative retraining process and, using this regularization term, directly produces a sparse model during the training process. Furthermore, our approach is also much easier and simpler to implement than existing methods. Experimental results on MNIST and CIFAR-10 show that our approach significantly reduces the number of filters of classical models such as *LeNet* and *VGG* while reaching the same or even better accuracy than the baseline models. Moreover, the trade-off between the sparsity and the accuracy is compared to other loss regularization terms based on the l_1 or l_2 norm as well as the SSL [1], NISP [2] and GAL [3] methods and shows that our approach is outperforming them.

Keywords: Deep Learning · Compression · Neural Networks · Architecture.

1 Introduction

Since the advent of *Deep Neural Networks* (DNNs) and especially *Deep Convolutional Neural Networks* (DCNNs) and their massively parallelized implementations [4][5], deep learning based methods have achieved state-of-the-art

* Corresponding author : anthony.berthelie@etu.uca.fr. This work has been sponsored by the Auvergne Regional Council and the European funds of regional development (FEDER)

performance in numerous visual tasks such as face recognition, semantic segmentation, object classification and detection, etc. [4][6][7][8][9]. Accompanied with the high performance, also high computation capabilities and large memory resources are needed as these models usually contain millions of parameters. These issues prevent them from running on resource-limited devices such as smartphones or embedded devices. Network compression is a common approach in this context, i.e. to reduce the inherent redundancy in the parameters and thus in the computation.

Numerous methods have been developed to obtain compact DNNs. Since a large number of these networks are built upon convolutional layers and since the convolution operations are the most computationally demanding, we are focusing on the reduction of these layers. A simple reduction strategy consists in removing non-relevant filters using pruning methods. For example, Li *et al.* [10] proposed to remove filters that are identified as having a small effect on the output accuracy. Another approach by Luo *et al.* [11] is evaluating information at the filter level using statistical and optimization methods.

Our approach is motivated and inspired by (1) previous works demonstrating the redundancy among the weights of a DCNN [12]; (2) numerous sparsity methods proposed in the literature [13] and (3) the fact that these sparsity methods have rarely been used to remove unimportant weights during training [1]. We therefore propose a new strategy, based on l_1/l_2 -norm, to obtain a subset of kernels with all weights equal to zero (such as that the associated filters can be removed). The main idea is to express the filter reduction problem by introducing sparsity on a set of pseudo-norms computed on each kernel but not directly on the kernels actual values. Figure 1 illustrates the general idea of our method. Each kernel of the network is transformed to a single value using a pseudo-norm. All these values are concatenated into a global vector (its size being the number of filters) called kernel norm vector. Our global kernel-sparsity is defined by the sparsity on this vector and is estimated by a l_1/l_2 -norm ratio. Since a kernel with all weights equal to zero produces a pseudo-norm of zero, the number of filters can be reduced by enforcing sparsity on the kernel norm vector. In this paper, we propose the l_1/l_2 -norm for two reasons: (1) the so-called l_1/l_2 -norm is a simple group norm to implement and (2) the use of the l_1 -norm can increase the performance, interpretability and sparsity of a model [14] [15][16] combined with the l_2 -norm allows to converge to stable solution and maintain sparsity at a good level.

We propose a l_1/l_2 -norm computed on the global vector (vector of kernel pseudo-norms) such that adding this sparsity term to minimize to the global loss will reduce the number of (non-zero) filters of a DCNN. Compared to other approaches, our method presents several advantages:




Fig. 1: Visual representation of our method and the computation of the kernel norm vector using a pseudo-norm.

1. All steps are done during training, i.e. no additional fine-tuning operations are needed.
2. Our method being based on simple l_1 and l_2 norms, is straightforward to implement and compute compared to other methods that remove weights during training.
3. As we are keeping track of the evolution of the network at every step during training, it is possible to choose the best model based on a trade-off between compression and accuracy.

In the following, we will first present existing work related to network pruning and weight sparsity, in Section 2. In Section 3, we describe our l_1/l_2 pseudo-norm method. Finally, in Section 4, we show experimental results of our method with *LeNet* and *VGG* network architectures trained on the MNIST and CIFAR-10 datasets. We demonstrate that our method is able to significantly improve the sparsity among convolutional layers in these DCNNs without significant drops in accuracy.

2 Related Work

Many studies have been done on DNN compression. Knowledge distillation [17][18] tackles the problem of transferring the encoded information from a bigger

model into a smaller one. Lowering numerical precision is also an extensive field [19][20][21]. Many works, are focused on designing compressed and optimal models architectures. SqueezeNet [22] and MobileNets [23] both propose structures of convolutional layers to improve memory and computation time. Some Neural Architecture Search (NAS) [24][25][26] methods use reinforcement learning and genetic algorithms to search the best possible networks designs for a given task. Depending on the size of the search space, finding an optimized model with these methods can be enormously time-consuming. However, the most promising approaches try to reduce the model redundancy and among them: parameter quantization [27][28] and network pruning [29][30][31][29][10][11]. Our method can be classified in this last category.

2.1 Network Pruning

Pruning methods are aiming to remove unimportant parameters of a neural network. Han *et al.* [27][29] proposed to prune parameters of *AlexNet* and *VGG* with connection pruning by setting a threshold and removing any parameters under it. As opposed to our method, most of the reduction is done on fully connected layers and not on convolutional layers. However, compression of convolutional layers is essential nowadays as new DNNs are mostly DCNNs with fewer fully connected layers *e.g.*, only 3.99% parameters of *Resnet* [9]. Closer to our approach, structured pruning methods are removing directly structured parts *e.g.*, kernels or layers, to compress CNNs. Li *et al.* [10] used l_1 -norm to remove filters. He *et al.* [32] used a LASSO regression based channel selection to prune filters. Channel pruning methods are preferred on widely-used DCNNs. For example, the selection of unimportant feature maps can be done using l_1 -regularization [33].

These past few years, numerous networks compression algorithms using pruning methods and achieving state-of-the-art results have emerged. Yu *et al.* [2] proposed a neurons importance score propagation (NISP) method based on the response of the final layers to evaluate the pruning impact of the prior layers. Zhuang *et al.* [34] developed discrimination-aware losses in order to determine the most useful channels in intermediate layers. Some methods such as Filter Pruning Via Geometric Median (FPGM) [35] are not focused on pruning filters with less importance but only by evaluating their redundancy. Similarly, Lin *et al.* [3] tackled the problem of redundant structures by proposing a generative adversarial learning method (GAL) (not only to remove filters, but also branches and blocks).

Still, standard pruning methods usually construct non structured and irregular connectivity in a network, leading to irregular memory access. In most of these approaches, the DNN is trained first. Then each parameter is evaluated to understand if it brings information to the network. If not, the parameter is removed. Therefore, a fine-tuning needs to be performed afterwards to restore the model accuracy. These steps take time. Most of them are done offline and need costly reiterations of decomposing and fine-tuning to find an optimal weight approximation maintaining high accuracy and high compression rate. Unlike these

methods, our approach is able to directly increase the sparsity of the network during training, identifying which kernels to prune without any considerable extra computational overhead.

2.2 Weight Sparsity

An important factor for the compression of a model is its sparsity *i.e.* the number of parameters set to zero. However, this sparsity must be structured in order to be memory-efficient and time-efficient. Liu *et al.* [36] obtained a sparsity of 90% on *AlexNet* with only 2% accuracy loss using sparse decomposition and a sparse matrix multiplication algorithm. This method also employed group Lasso [37], an efficient regularization to learn sparse structures. It is also used by Wen *et al.* [1] to regularize the structure of a DNN at different levels (*i.e.* filters, channels, filter shapes and layer depth). This approach leads to DNNs with reduced computational cost and efficient acceleration due to the structured sparsity induced by the method. We propose to use a different type of regularization based on the norm ratio l_1/l_q [38][13]. It allows to dynamically maximize the sparsity of a model with one hyper-parameter (q) without additional iterations and severe drops in accuracy while being straightforward to implement.

3 Training with kernel-sparsity

We mainly focus on inducing sparsity on convolutional layers to regularize and compress the structure of DCNNs during the training steps. We propose a generic method to regularize DCNNs using the l_1/l_2 pseudo-norm.

3.1 Kernel-sparsity regularization

Let \mathcal{N} be a DCNN with L convolutional layers. We define $W^{l,k}$ as the $k^{th} \in \{1, \dots, N_k^l\}$ 3d-tensor (kernel) associated with the l^{th} convolutional layer. Thus, a weight of kernel k in the convolutional layer l is defined as: and $W_{w,h,c}^{l,k} \in \mathbb{R}^{N_w^l, N_h^l, N_c^l}$ the (width, height, channel) weight of kernel k of layer l .

$$W_{w,h,c}^{l,k} \in \mathbb{R}^{N_w^l, N_h^l, N_c^l} \quad (1)$$

Here, $w \in \{1, \dots, N_w^l\}$ is the column, $h \in \{1, \dots, N_h^l\}$ is the row and $c \in \{1, \dots, N_c^l\}$ is the channel index of the k^{th} kernel matrix in the convolutional layer l . The key idea is to express sparsity on pseudo-norms of kernels. Let n_k^l be the pseudo-norm defined by the l_1 -norm of the flattened kernel $W^{l,k}$:

$$n_k^l \doteq \sum_{w=1}^{N_w^l} \sum_{h=1}^{N_h^l} \sum_{c=1}^{N_c^l} \frac{|W_{w,h,c}^{l,k}|}{N_k^l} \quad (2)$$

The vector \vec{N}^l concatenates, for layer l , the N_k^l norms n_k^l :

$$\vec{N}^l \doteq \left\| \left\| n_k^l \right\|_{k=1}^{N_k^l} \right\| \quad (3)$$

We introduce kernel-sparsity for a layer as a value linked to the number of kernels of this layer with all weights equal to zeros. Therefore, the kernel-sparsity of layer l can be linked with the number of values of the vector \vec{N}^l equal to zero. Global kernel-sparsity can be expressed from the concatenation of vectors \vec{N}^l for each layer:

$$\vec{N} \doteq \left\| \left\| \vec{N}^l \right\|_{l=1}^L \right\| \quad (4)$$

For better understanding, we visualize these operations in Figure 2. In order to normalize the value of N , each of its component is divided by the number of values (or norms) that it contains. Finally, the global kernel-sparsity is defined by the sparsity of \vec{N} and can be estimated by a l_1/l_2 ratio function:

$$\mathcal{L}_s \doteq \frac{\vec{N}_1}{\vec{N}_2} \quad (5)$$

Minimizing this term will encourage zero-valued coefficients (numerator), corresponding to the different kernels, while keeping the remaining coefficients at large values (denominator), thus producing convolution layers with few non-zero kernels.

3.2 Training with kernel-sparsity regularization

Let $\mathcal{L}_{\mathcal{N}}$ be the loss function that is minimized to find the optimal weight configuration for a given task (e.g. cross entropy). We propose to simply add the




Fig. 2: Visualization of the computation of the kernels pseudo-norm and how the global kernel norm vector \vec{N} is obtained.

kernel-sparsity regularization term weighted by the coefficient $\lambda \in \mathbb{R}$:

$$\mathcal{L}_{all} = \mathcal{L}_{\mathcal{N}} + \lambda \mathcal{L}_s . \quad (6)$$

We will discuss how to set an appropriate values of λ in the experimental section.

3.3 Setting kernels to zero

Our method induces sparsity in a DCNN, i.e. the pseudo-norm regularization pushes some kernels to have only zero-valued coefficients. However, in practice, during optimization, the actual values of these kernels will not be exactly zero but very small. Thus, to compress the network effectively, our approach identifies these kernels during training and forces them to be zero in order to remove them.

More specifically, the algorithm works as follows: each pseudo-norm of the kernels is contained in the global kernel pseudo-norm vector \vec{N} . Thus, at each epoch, we normalize the values of \vec{N} so that $\sum_{i=1}^K \vec{N}_i = 1$. Sorting these vectors in ascending order will allow us to objectively determine which pseudo-norms are the smallest. We then define a percentage (or a threshold) under which the cumulative sum of these sorted values is judged too small to be kept, i.e. the corresponding filters are considered unimportant and set to zero. Once the weights of a kernel are set to zero, they are keeping this value until the end of the training, and these parameters are no longer updated. This ensures that the potential errors and imprecision introduced by removing these kernels can be compensated by the remaining kernels during the training converging to a stable solution with high accuracy.

To summarize, our approach consists of two steps at each epoch:

1. The l_1/l_2 pseudo-norm is computed on each kernel of the model and is integrated to the loss function. Thus the training stage is minimizing the loss function and inducing sparsity at the kernel level, pushing some weights to have a near zero value.
2. Sort kernels according to their ascending normalized pseudo-norm and compute a cumulative sum vector from the sorted normalized pseudo-norm vector. The kernels participating to the cumulative sum under a threshold t are removed. This set of operations aims at keeping kernels that produce more than $t\%$ of the global norm.

4 Experiments

We evaluated the performance of the l_1/l_2 pseudo-norm on two classification models (*LeNet* and *VGG*) and two datasets: MNIST and CIFAR-10. Our method is implemented in *Pytorch*, running on various Nvidia GPUs using CUDA. The weights of the networks are initialized randomly and hyper-parameters are selected manually for optimal results. The chosen λ value is the one allowing the model to have an accuracy close to its baseline accuracy while sparsifying the





| EPOCH | KERNELS | \mathcal{L}_s |
|-------|--|-----------------|
| 0 |  | 0.0226 |
| 4 |  | 0.0220 |
| 20 |  | 0.0189 |
| 130 |  | 0.0118 |

Fig. 3: Evolution of the 20 kernels of the first convolutional layer of *LeNet* and the global kernel-sparsity regularization term \mathcal{L}_s during a training on the MNIST dataset. For a better visualization, each kernel is flattened from 3 dimensions to 2 dimensions.

most the kernels norms. In all the experiments, the threshold under which the kernels are removed by evaluating the cumulative sum of the smallest norms is set to 1%. We found that this value was the best trade-off between a converging accuracy of the models and a slow removal of the kernels during the training phase.

4.1 Experiments on *LeNet*

In the experiments with *LeNet* [39], we investigate the effectiveness of the l_1/l_2 pseudo-norm on the MNIST and CIFAR-10 datasets. In order to compare our results with state-of-the-art methods such as SSL [1], NISP [2] and GAL [3], we decide to chose the *LeNet* model implemented by *Caffe*. All these methods are using evaluation and regression at different level i.e Lasso-group regression at different level of a convolutional layer in the SSL method, which makes it more complex to implement than our method. There is no data augmentation for the training on both datasets.

LeNet on MNIST: As previously described, the l_1/l_2 pseudo-norm is applied on the filters of a DCNN to penalize them. Hence our method is inducing sparsity among the filters of the convolutional layers in *LeNet*. To visualize the effect on our approach on the kernels, Figure 3 shows the evolution of the kernels of the first convolutional layer of *LeNet* during a training on MNIST with our kernel-sparsity regularization. We see that the kernel-sparsity term \mathcal{L}_s is decreasing epoch after epoch and that the number of filters in the layer is also decreasing with it. The complete evolution of the kernel-sparsity \mathcal{L}_s can be seen on Figure 4. \mathcal{L}_s being computed on the pseudo-norm of the kernels and kernels weights being set to zero over time: this result shows the effectiveness of our method.

Table 1 summarizes the results on MNIST of different methods. In the best case scenario that we have tested, the l_1/l_2 pseudo-norm with a λ value set to 0.5 is able to achieve an accuracy better than the baseline by 0.2%. Furthermore the number of filters is dropping drastically in both convolutional layers respectively




Fig. 4: Evolution of the cross-entropy loss function \mathcal{L}_N and the global kernel-sparsity regularization term \mathcal{L}_s during training with $\lambda = 0.5$. Evaluations are done on *LeNet* on the MNIST dataset. Vertical lines show which epochs were taken to construct Figure 3.

from 20 to 5 and from 50 to 18. Compared to the other state-of-the-art methods, the l_1/l_2 pseudo-norm is able to achieve a better accuracy while penalizing more filters too. We also compare our method to the l_1 -norm and l_2 -norm. During our evaluations, both of these norms were able to reach a higher level of sparsity by setting to zero more kernels. However they were never able to reach the same or a better level of accuracy than the baseline.

To visualize the effect of our method on the parameters, we show the learned filters of the first convolutional layer in Figure 5. For $\lambda = 0.5$ and for different level of sparsity, it can be seen that the number of remaining filters can be set to only 2 or 4. Furthermore between the baseline and our method, the accuracy is the same or is increased. This shows that there is effectively a large amount of redundancy between filters and that most of them are not required. Moreover, compared to the baseline, it seems that the remaining filters are more structured, with more regular patterns. This assumption seems especially true when only two filters are remaining. Thus, we arrive at the same conclusion than [1]: the




Fig. 5: Learned filters of the first convolutional of *LeNet* on MNIST. Top is *LeNet* baseline, middle and bottom are l_1/l_2 -norm with $\lambda = 0.5$ and different level of sparsity.

| Method | λ | Error | Conv1 Filter # (Sparsity) | Conv2 filter # (Sparsity) | Total Sparsity |
|-----------|-----------|-------------|------------------------------|------------------------------|-------------------|
| Baseline | - | 0.9% | 20 | 50 | 0% |
| l_1 | 0.5 | 1.2% | 4 (80%) | 5 (90%) | 87.1% |
| l_2 | 0.5 | 1.2% | 3 (85%) | 5 (90%) | 88.6% |
| SSL 1 | - | 0.8% | 5 (75%) | 19 (62%) | 65.7% |
| SSL 2 | - | 1.0% | 3 (85%) | 12 (76%) | 78.6% |
| NISP | - | 0.8% | 10 (50%) | 25 (50%) | 50.0% |
| GAL | - | 1.0% | 2 (90%) | 15 (70%) | 75.7% |
| l_1/l_2 | 0.5 | 0.7% | 5 (75%) | 18 (64%) | 67.1% |

Table 1: Results after penalizing unimportant filters in *LeNet* on MNIST. Baseline is the simple *LeNet* Caffe model. l_1 and l_2 are the best results found by using the l_1 -norm and l_2 -norm regularization on the kernels. SSL, NISP and GAL are the pruning methods respectively from [1], [2] and [3]. l_1/l_2 is our method with $\lambda = 0.5$.

baseline has a high freedom in the parameter space and our method is able to obtain the same accuracy by optimizing the filters into more regularized patterns.

LeNet on CIFAR-10: In order to test the l_1/l_2 pseudo-norm and visualize its effect on a more difficult classification task than MNIST, we decided to use the CIFAR-10 dataset with the same *LeNet* model. The results are summarized in Table 2. The baseline *LeNet* is not performing as well on CIFAR-10 than it is performing on MNIST, i.e. the classification accuracy is only around 70%. As a result, the accuracy of our approach also drops but the l_1/l_2 pseudo-norm is still able to perform well on this model, even for this classification task. With a λ value set to 0.7, we are able to decrease the number of filters in the first and the second convolutional layers respectively from 20 to 10 and from 50 to 25, which means that half of the filters of *LeNet* are removed. With this configuration, our method performs 1.7% worse than the baseline. We were able to remove up to 80% of the filters in our experiments, but the resulting accuracy was too low to be interesting (more than 20% behind the baseline). Hence, more filters are needed in order to classify correctly the CIFAR-10 dataset compared to the MNIST dataset. The best trade-off between filters and accuracy that we




Fig. 6: Learned filters of the first convolutional layer of *LeNet* on CIFAR-10. Top is *LeNet* baseline, middle and bottom are l_1/l_2 -norm with $\lambda = 0.7$ and different level of sparsity.

| Method | λ | Error | Conv1 Filter # (Sparsity) | Conv2 filter # (Sparsity) | Total Sparsity |
|-----------|-----------|--------------|------------------------------|------------------------------|-------------------|
| Baseline | - | 28.4% | 20 | 50 | 0% |
| l_1 | 0.7 | 35.4% | 7 (65%) | 14 (72%) | 70.0% |
| l_2 | 0.7 | 29.8% | 12 (40%) | 24 (52%) | 48.6% |
| l_1/l_2 | 0.7 | 30.1% | 10 (50%) | 25 (50%) | 50.0% |
| l_1/l_2 | 0.7 | 29.3% | 14 (30%) | 30 (40%) | 37.1% |

Table 2: Results after penalizing unimportant filters in *LeNet* on CIFAR-10. Baseline is the simple *LeNet* Caffe model. l_1/l_2 is our method with different coefficient of regulation $\lambda = 0.7$.

found was still with a value of λ set to 0.7. In both convolutional layers, the number of filters is dropping respectively from 20 to 14 and from 50 to 30. This means that our method is able to zero out more than a third of the filters with only a drop of 0.9% in accuracy. Compared to the l_1 -norm and the l_2 -norm, the l_1/l_2 pseudo-norm also shows good results. Indeed, both the l_1 -norm and l_2 -norm were unable to reach the same level of accuracy and setting to zero as many filters as the l_1/l_2 pseudo-norm can do.

As previously done with the MNIST dataset, we visualize the learned filters of the first convolutional layer in Figure 6. From this visualization, we can draw the same conclusion than with the MNIST dataset. The more we are removing filters, the more the remaining ones seem to have a defined structure, as opposed to the baseline where each of the filters seems blurry. It is even more remarkable when we let our algorithm run until only a couple of filters are remaining. Even if the model does not reach a satisfactory accuracy, the two remaining filters have learned remarkable patterns. Thus, the l_1/l_2 pseudo-norm is still able to smooth a high freedom of parameter space into fewer filters with more regularized patterns.

4.2 VGG on CIFAR10

To demonstrate the generalization of our method on larger DNNs, we evaluate the performance of our method on the well-known *VGG* [6], a deeper model than *LeNet*, with several convolutional layers. A *VGG* model can have different sizes, notably depending on the number of layers. We chose the *VGG11* model with a total of 8 convolutional layers. We implemented it using *Pytorch*, running on various Nvidia GPUs using CUDA. The model is trained without data augmentation and evaluated on the CIFAR-10 dataset. In this experiment, the kernels pseudo-norms are not normalized on the full network, which explains why the λ values are smaller than the ones used with *LeNet*.

With *LeNet*, the l_1/l_2 pseudo-norm method was applied on only 2 convolutional layers, with 50 filters at most in the second convolutional layer. In *VGG11* our method is applied on 8 different convolutional layers with a number of filters set to 64 in the first convolutional layer and a maximum of 512 filters in the last four convolutional layers. The results are shown in Table 3. The baseline model, with all the filters and a classical loss function (cross-entropy), obtains an error of 17.6% on the test dataset. Using the l_1/l_2 pseudo-norm with a λ set to 0.005, the model achieves a classification accuracy roughly 1% inferior to the baseline. However the number of filters is vastly reduced. Moreover, it seems that the deeper we go in the network, the more the proportion of filter sets to zero is important. For example, the second convolutional layer has around 10% of its filters set to zero while the last convolutional layer has over 65% of filters set to zero. Thus we could deduce that the last convolutional layers keep less important information for the model than the first ones or that there is more redundancy in the last layers. However, the first convolutional layer seems to be an exception as approximately half of its filters can be removed. We suppose that the shapes learned in the first layer are not decisive for the model and can be balanced by the following layers and the more defined shapes that they have assimilated.


By decreasing the λ coefficient to 0.001, we confirm the results that the last convolutional layers seem to contain more filters with non decisive or redundant information than the first ones. Indeed, only the last two layers have filters set

| Method | λ | Error | Conv1 to conv 8 Filter # | Total Sparsity |
|-----------|-----------|--------------|--|----------------|
| Baseline | - | 17.6% | 64 - 128 - 256 - 256 - 512 - 512 - 512 - 512 | 0% |
| l_1 | 0.0001 | 19.3% | 52 - 128 - 255 - 256 - 175 - 147 - 97 - 123 | 55.2% |
| l_2 | 0.005 | 18.2% | 64 - 128 - 256 - 256 - 511 - 474 - 434 - 299 | 12% |
| l_1/l_2 | 0.005 | 18.8% | 35 - 115 - 238 - 176 - 354 - 195 - 190 - 175 | 46.3% |
| l_1/l_2 | 0.001 | 16.8% | 64 - 128 - 256 - 256 - 512 - 512 - 510 - 380 | 5% |

Table 3: Results after penalizing unimportant filters in *VGG11* on CIFAR-10. Baseline is the *VGG11* network baseline. l_1/l_2 is our method with different coefficient of regulation λ .

to zero. But more importantly, the removal of a few filters in the last two convolutional layers leads to a classification error of only 16.8%, which is 0.8% less than the baseline. Thus, our method, by only removing a few filters, is able to achieve a better accuracy than the baseline model. Compared to the l_1 -norm and the l_2 -norm, the l_1/l_2 pseudo-norm is also performing well. The l_1 -norm is able to zero out numerous filters but is unable to achieve a correct level of accuracy, always performing worse than the baseline or our approach. Nearly the same conclusions can be drawn from the l_2 -norm. Under certain conditions, the l_2 -norm is able to zero out slightly more filters than our method in the last convolution layers. However, the models are not able to obtain a satisfying accuracy, always around 1% behind the baseline.

In order to conclude this study, we visualize in Figure 7 the evolution of the accuracy of the model against the number of kernel set to zero in the first and second convolutional layers for different coefficient of regularization λ . When $\lambda = 0$, the l_1/l_2 pseudo-norm is not taken into account, resulting into the baseline model. The order that the filters are set to zero is determined by the filters pseudo-norm arranged by ascending order. These tests are done at a single convolutional layer level. Meaning that during training, the only filters that are evaluated and set to zero are the ones belonging to the studied layer. The other layers are remaining untouched. We visualize that for both layers, we are able to set numerous filters to zero without a noticeable decrease of the accuracy, even when our method is not active. This result shows that there is unimportant information in the layer and that it is possible to remove it, even if there are no methods that are defined to emphasize this phenomenon. With the implementation of the l_1/l_2 pseudo-norm ($\lambda > 0$), we see that (1) more kernels are set to zero before the beginning of the accuracy drop compared to the baseline and (2) a greater λ value means that more kernels are zeroed-out but at the price of an inferior accuracy. Based on these conclusions, our method is increasing the sparsity of the filters within a layer, shifting information between them in order to centralize the information. However this sparsity has its limits. The more we force it (with a significant λ value), the more we increase the chances to lose important information that could be never recovered.



(a) First convolutional layer of *VGG11* (b) Second convolutional layer of *VGG11*

Fig. 7: Visualization of the effect of setting kernels to zero in the first two convolutional layers of *VGG11* against the accuracy of the network. Each line represents a different value for the coefficient of regulation λ of the l_1/l_2 pseudo-norm method.

5 Conclusion

In this work, we proposed a new regularization approach for inducing kernel sparsity in DCNNs based on the l_1/l_2 pseudo-norm. This method reorganizes the weights of the convolutional layers in order to learn more compact structures during training. These compact DCNNs can reach almost the same accuracy as the original models and in some cases even perform better. Our experiments have demonstrated the benefits of our approach and its generalization to deep structures: it is straightforward to implement, it operates during the training process and it is possible to choose between the compactness and the accuracy of the model. So far, we have only applied our method to classification problems. To go beyond, in the future, the method needs to be applied on deeper models such as Resnet [9] and bigger datasets. Autoencoders, fully convolutional networks and segmentation problems are also an important focus. Furthermore, it would be interesting to generalize the l_1/l_2 -norm to the l_1/l_q -norm to study its properties in more detail and improve on different model structures and learning problems.

References

1. W. Wen, C. Wu, Y. Wang, Y. Chen, and H. Li, "Learning structured sparsity in deep neural networks," *NIPS*, 2016.
2. R. Yu, A. Li, C. Chen, J. Lai, V. I. Morariu, X. Han, M. Gao, C. Lin, and L. S. Davis, "NISP: pruning networks using neuron importance score propagation," *CVPR*, 2018.
3. S. Lin, R. Ji, C. Yan, B. Zhang, L. Cao, Q. Ye, F. Huang, and D. S. Doermann, "Towards optimal structured CNN pruning via generative adversarial learning," *CVPR*, 2019.
4. A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," *NIPS*, 2012.
5. Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, 2015.
6. K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale Image recognition," *CoRR*, 2015.
7. C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich, "Going deeper with convolutions," *CVPR*, 2015.
8. K. He and J. Sun, "Convolutional neural networks at constrained time cost", 2015.
9. K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," *CVPR*, 2016.
10. H. Li, A. Kadav, I. Durdanovic, H. Samet, and H. P. Graf, "Pruning filters for efficient ConvNets," *ICLR*, 2017.
11. J.-H. Luo, J. Wu, and W. Lin, "Thinet: A filter level pruning method for deep neural network compression," *ICCV*, 2017.
12. M. Jaderberg, A. Vedaldi, and A. Zisserman, "Speeding up convolutional neural networks with low rank expansions," *BMVC*, 2014.
13. F. Bach, R. Jenatton, J. Mairal, and G. Obozinski, "Optimization with sparsity-inducing penalties," *Found. Trends Mach. Learn.*, 2012.
14. J. Huang and T. Zhang, "The benefit of group sparsity," *Ann. Statist.*, 2010.
15. B. Turlach, W. Venables, and S. Wright, "Simultaneous variable selection," *Technometrics*, 2000.

16. M. Yuan and Y. Lin, "Model selection and estimation in regression with grouped variables," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 2006.
17. Y. N. Dauphin and Y. Bengio, "Big neural networks waste capacity," 2013.
18. J. Ba and R. Caruana, "Do deep nets really need to be deep?" *NIPS*, 2014.
19. S. Gupta, A. Agrawal, K. Gopalakrishnan, and P. Narayanan, "Deep Learning with Limited Numerical Precision," *ICML*, 2015.
20. M. Courbariaux, Y. Bengio, and J.-P. David, "Training deep neural networks with low precision multiplications," *ICLR*, 2014.
21. D. Williamson, "Dynamically scaled fixed point arithmetic," *IEEE Pacific Rim Conference on Communications, Computers and Signal Processing Conference Proceedings*, 1991.
22. F. N. Iandola, S. Han, M. W. Moskewicz, K. Ashraf, W. J. Dally, and K. Keutzer, "Squeezenet: Alexnet-level accuracy with 50x fewer parameters and 0.5 mb model size," 2016.
23. A. G. Howard, M. Zhu, B. Chen, D. Kalenichenko, W. Wang, T. Weyand, M. Andreetto, and H. Adam, "Mobilenets: Efficient convolutional neural networks for mobile vision applications," 2017.
24. R. Miikkulainen, J. Liang, E. Meyerson, A. Rawal, D. Fink, O. Francon, B. Raju, H. Shahrzad, A. Navruzyan, N. Duffy, and B. Hodjat, "Evolving Deep Neural Networks," 2017.
25. M. Tan, B. Chen, R. Pang, V. Vasudevan, and Q. V. Le, "MnasNet: Platform-Aware Neural Architecture Search for Mobile," *CVPR*, 2019.
26. Y. He, J. Lin, Z. Liu, H. Wang, L.-J. Li, and S. Han, "Amc: Automl for model compression and acceleration on mobile devices," *ECCV*, 2018.
27. S. Han, H. Mao, and W. J. Dally, "Deep Compression - Compressing Deep Neural Networks with Pruning, Trained Quantization and Huffman Coding," *ICLR*, 2016.
28. Y. Choi, M. El-Khamy, and J. Lee, "Towards the limit of network quantization," *ICLR*, 2017.
29. S. Han, J. Pool, J. Tran, and W. Dally, "Learning both weights and connections for efficient neural network," *NIPS*, 2015.
30. S. Anwar, K. Hwang, and W. Sung, "Structured pruning of deep convolutional neural networks," *ACM Journal on Emerging Technologies in Computing Systems*, 2017.
31. P. Molchanov, S. Tyree, T. Karras, T. Aila, and J. Kautz, "Pruning convolutional neural networks for resource efficient transfer learning," *ICLR*, 2017.
32. Y. He, X. Zhang, and J. Sun, "Channel pruning for accelerating very deep neural networks," *ICCV*, 2017.
33. Z. Liu, J. Li, Z. Shen, G. Huang, S. Yan, and C. Zhang, "Learning Efficient Convolutional Networks through Network Slimming," *ICCV*, 2017.
34. Z. Zhuang, M. Tan, B. Zhuang, J. Liu, Y. Guo, Q. Wu, J. Huang, and J. Zhu, "Discrimination-aware channel pruning for deep neural networks," in *NIPS*, 2018.
35. Y. He, P. Liu, Z. Wang, Z. Hu, and Y. Yang, "Filter pruning via geometric median for deep convolutional neural networks acceleration," *CVPR*, 2019.
36. Baoyuan Liu, Min Wang, H. Foroosh, M. Tappen, and M. Pensky, "Sparse convolutional neural networks," *CVPR*, 2015.
37. M. Yuan and Y. Lin, "Model selection and estimation in regression with grouped variables," *Journal of the Royal Statistical Society Series B*, 2006.
38. J. Liu and J. Ye, "Efficient l_1/l_q norm regularization," 2010.
39. Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proceedings of the IEEE*, 1998.