

## Neural network relief: a pruning algorithm based on neural activity

Dekhovich, Aleksandr; Tax, David M.J.; Sluiter, Marcel H.F.; Bessa, Miguel A.

**DOI**

[10.1007/s10994-024-06516-z](https://doi.org/10.1007/s10994-024-06516-z)

**Publication date**

2024

**Document Version**

Final published version

**Published in**

Machine Learning

**Citation (APA)**

Dekhovich, A., Tax, D. M. J., Sluiter, M. H. F., & Bessa, M. A. (2024). Neural network relief: a pruning algorithm based on neural activity. *Machine Learning*, 113(5), 2597-2618. <https://doi.org/10.1007/s10994-024-06516-z>

**Important note**

To cite this publication, please use the final published version (if applicable). Please check the document version above.

**Copyright**

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

**Takedown policy**

Please contact us and provide details if you believe this document breaches copyrights. We will remove access to the work immediately and investigate your claim.

***Green Open Access added to TU Delft Institutional Repository***


***'You share, we take care!' - Taverne project***

**<https://www.openaccess.nl/en/you-share-we-take-care>**

Otherwise as indicated in the copyright section: the publisher is the copyright holder of this work and the author uses the Dutch legislation to make this work public.



# Neural network relief: a pruning algorithm based on neural activity

Aleksandr Dekhovich<sup>1</sup> · David M. J. Tax<sup>2</sup> · Marcel H. F. Sluiter<sup>1</sup> · Miguel A. Bessa<sup>3</sup> 

Received: 19 December 2022 / Revised: 15 November 2023 / Accepted: 18 January 2024

© The Author(s), under exclusive licence to Springer Science+Business Media LLC, part of Springer Nature 2024

## Abstract

Current deep neural networks (DNNs) are overparameterized and use most of their neuronal connections during inference for each task. The human brain, however, developed specialized regions for different tasks and performs inference with a small fraction of its neuronal connections. We propose an iterative pruning strategy introducing a simple importance-score metric that deactivates unimportant connections, tackling overparameterization in DNNs and modulating the firing patterns. The aim is to find the smallest number of connections that is still capable of solving a given task with comparable accuracy, i.e. a simpler subnetwork. We achieve comparable performance for LeNet architectures on MNIST, and significantly higher parameter compression than state-of-the-art algorithms for VGG and ResNet architectures on CIFAR-10/100 and Tiny-ImageNet. Our approach also performs well for the two different optimizers considered—Adam and SGD. The algorithm is not designed to minimize FLOPs when considering current hardware and software implementations, although it performs reasonably when compared to the state of the art.

**Keywords** Neural network pruning · Iterative pruning · Connection pruning · Sparsity

## 1 Introduction

Pruning is a common technique for neural network compression (Karnin, 1990), where the main goal is to reduce memory and computational costs of inference. Pruning assumes particular relevance for deep neural networks because modern architectures involve several millions of parameters. Existing pruning methods are based on different strategies,

---

Editor: Andrea Passerini.

---

✉ Miguel A. Bessa  
miguel\_bessa@brown.edu

<sup>1</sup> Department of Materials Science and Engineering, Delft University of Technology, Mekelweg 2, 2628 CD Delft, The Netherlands

<sup>2</sup> Pattern Recognition and Bioinformatics Laboratory, Delft University of Technology, Van Mourik Broekmanweg 6, 2628 XE Delft, The Netherlands

<sup>3</sup> School of Engineering, Brown University, 184 Hope St., Providence, RI 02912, USA

e.g. Hessian analysis (LeCun et al., 1990; Hassibi and Stork, 1993), magnitudes of weights (Han et al., 2015), data-driven approaches (Hu et al., 2016; Ancona et al., 2020), among others (Ullrich et al., 2017; Molchanov et al., 2017). Pruning can be done in one shot (Zhao et al., 2019) or in an iterative way (Frankle and Carbin, 2019), and it is possible to prune connections (LeCun et al., 1990; Hassibi and Stork, 1993; Dong et al., 2017; Han et al., 2016), neurons (Molchanov et al., 2019; Louizos et al., 2017; Ancona et al., 2020) or filters for convolutional layers (He et al., 2017, 2018). The typical pruning pipeline includes three stages: training the original network, pruning parameters and fine-tuning. Recently, interesting solutions for the third stage have been suggested that involve weight rewinding (Frankle and Carbin, 2019) and learning rate rewinding (Renda et al., 2020).

We have developed an algorithm that aims to prune a network without a significant decrease in accuracy after every iteration by keeping the signal in the network at some pre-defined level close to the original one. This contrasts with strategies that use a pre-defined ratio of parameters to prune (Hu et al., 2016; Ancona et al., 2020; Frankle and Carbin, 2019) which may lead to a drastic drop in accuracy for relatively high ratios. We look at the local behaviour of a particular connection and its contribution to the neuron, but not at the output or the loss function. Our aim is to deactivate unimportant connections for a given problem in order to free them for other tasks—a crucial step towards novel architectural continual learning strategies (Mallya and Lazebnik, 2018; Mallya et al., 2018). Sparse architectures are also more robust to noisy data (Ahmad and Scheinkman, 2019) and exhibit benefits in the context of adversarial training (Ye et al., 2019; Liao et al., 2022). Therefore, we focus on reducing the number of parameters of deep learning models targeting these scenarios, although we expect a reduction of the computational complexity (FLOPs) as well.

Our iterative pruning algorithm is based on an importance score metric proposed herein that quantifies the relevance of each connection to the local neuron behaviour. We show compressions of more than 50 times for VGG (Simonyan and Zisserman, 2015) architectures on CIFAR-10 and Tiny-ImageNet datasets with a marginal drop of accuracy. We also apply our method to ResNets (He et al., 2016), achieving better parameter compression than state-of-the-art algorithms with a comparable decrease of accuracy on CIFAR-10 dataset. In addition, we visualise the effects of our pruning strategy on the information propagation through the network, and we observe a significant homogenization of the importance of the pruned neurons. We associate this homogenization with the notion of neural network relief: using fewer neuronal connections and distributing importance among them.

## 2 Related work

One of the first works eliminating unimportant connections in relatively small networks proposed analyzing the Hessian of the loss function (LeCun et al., 1990; Hassibi and Stork, 1993) without network retraining. This idea was further developed for convolutional networks (Dong et al., 2017; Lebedev and Lempitsky, 2016). However, computing second-order derivatives by calculating the Hessian is costly.

The magnitude-based approach (Han et al., 2015) is simple and fast because it only involves pruning the weights with the smallest magnitude. This assumes that parameters with small magnitudes do not contribute significantly to the resulting performance. For convolutional layers, the sum of kernels' elements in the filters is considered and filters with the smallest sum are pruned (Li et al., 2017). Iterative magnitude pruning is applied

for finding winning tickets (Frankle and Carbin, 2019)—the minimal subnetwork that can be trained at least as well as the original one with the same hyperparameters.

Other algorithms use input data to reduce the number of parameters. In Hu et al. (2016) pruning depends on the ratio of zero activations using ReLU function. So, the neurons that do not fire frequently enough are eliminated. A greedy approach is used for ThiNet (Luo et al., 2017), where channels that do not affect the resulting sum in convolutional layers are removed by solving an optimization problem. A similar algorithm is used in He et al. (2017), but the optimization problem is solved with LASSO regression for determining ineffective channels. The property of convexity and sparsity that ReLU produces provides analytical boundaries for Net-trim (Aghasi et al., 2017) by solving a convex optimization problem. The game-theoretic approach with Shapley values (Ancona et al., 2020) demonstrates good performance in a low-data regime, i.e. one-shot pruning without retraining, where neurons within one layer are considered as players in a cooperative game.

NISP (Yu et al., 2018) finds the contribution of neurons to the last layer before classification, while iSparse framework (Garg and Candan, 2020) trains sparse networks eliminating the connections that do not contribute to the output. A pruning strategy where the connections' contribution is based on computing the loss function derivatives was also suggested in Lee et al. (2019).

Additional methods of interest include Bayesian weight pruning (Molchanov et al., 2017) and Bayesian compression (Louizos et al., 2017) to prune neurons and aim at computational efficiency. SSS (Huang and Wang, 2018) introduces a scaling factor and sparsity constraints on this factor to scale neurons' or blocks' outputs. Similarly, SNLI (Ye et al., 2018) uses ISTA (Beck and Teboulle, 2009) to update the scaling parameter in Batch Normalization of convolutional layers to obtain a sparse representation. HRank (Lin et al., 2020) calculates the average rank of feature maps and prunes filters with the lowest ones. FSABP (Geng and Niu, 2022) finds filters that extract similar information and prunes them. As a result, filters that provide diversity in feature maps are retained. In the same way, CHIP (Sui et al., 2021) measures correlations between feature maps truncating feature maps with the lowest independence scores.

### 3 Our pruning approach

Our goal is to eliminate connections in a layer that, on average, provide a weak contribution to the next layer. We believe that a low magnitude of a weight does not mean that information passing through the connection has to have a negligible contribution. The converse is also assumed to be true: if a high weight magnitude is multiplied by a weak signal from the neuron (or even zero) then the contribution of that product is relatively insignificant in comparison with other input signals in the neuron, as long as this holds for most data. Therefore, our strategy contrasts significantly with magnitude-based pruning (Han et al., 2015).

#### 3.1 Fully connected layers

Assume that we have a pruning set  $\mathbf{X}^{(l-1)} = \{\mathbf{x}_1^{(l-1)}, \dots, \mathbf{x}_N^{(l-1)}\}$  with  $N$  samples, where each datapoint  $\mathbf{x}_n^{(l-1)} = (x_{n1}^{(l-1)}, \dots, x_{nm_{l-1}}^{(l-1)}) \in \mathbb{R}^{m_{l-1}}$  is the input for layer  $l - 1$  with dimension  $m_{l-1}$

and where  $1 \leq l \leq L$ . We define the *importance* of the connection between neuron  $i$  of layer  $l - 1$  and neuron  $j$  of layer  $l$  as:

$$s_{ij}^{(l)} = \frac{\overline{|w_{ij}^{(l)} x_i^{(l-1)}|}}{\sum_{k=1}^{m_{l-1}} \overline{|w_{kj}^{(l)} x_k^{(l-1)}|} + |b_j^{(l)}|}, \quad (1)$$

where  $\overline{|w_{ij}^{(l)} x_i^{(l-1)}|} = \frac{1}{N} \sum_{n=1}^N |w_{ij}^{(l)} x_{ni}^{(l-1)}|$ , and  $w_{ij}^{(l)}$  is the corresponding weight between neurons  $i$  and  $j$  (see Fig. 1) and  $b_j^{(l)}$  is the bias associated to neuron  $j$ . The importance score for the bias of neuron  $j$  is  $s_{m_{l-1}+1j}^{(l)} = \frac{|b_j^{(l)}|}{\sum_{k=1}^{m_{l-1}} \overline{|w_{kj}^{(l)} x_k^{(l-1)}|} + |b_j^{(l)}|}$ . The denominator corresponds to the total importance in the neuron  $j$  of layer  $l$  that we denote as  $S_j^{(l)} = \sum_{k=1}^{m_{l-1}} \overline{|w_{kj}^{(l)} x_k^{(l-1)}|} + |b_j^{(l)}|$ ,  $1 \leq j \leq m_l$ .

**Algorithm 1** Fully connected layers pruning

---


```

1: function FC_PRUNING(network, X,  $\alpha$ )
2:    $\mathbf{X}^{(0)} \leftarrow \mathbf{X}$ 
3:   for every fc_layer  $l$  in FC_Layers do
4:      $\mathbf{X}^{(l)} \leftarrow \text{fc\_layer}(\mathbf{X}^{(l-1)})$ 
5:     for every neuron  $j$  in fc_layer  $l$  do
6:       compute importance scores  $s_{ij}^{(l)}$  for every incoming connection  $i$ 
       using (1).
7:        $\hat{s}_{ij}^{(l)} = \text{Sort}(s_{ij}^{(l)}, \text{order} = \text{descending})$ 
8:        $p_0 = \min\{p : \sum_{i=1}^p \hat{s}_{ij}^{(l)} \geq \alpha\}$ 
9:       prune connections with importance score  $s_{ij}^{(l)} < \hat{s}_{p_0j}^{(l)}$ 
10:    end for
11:  end for
12:  return pruned network
13: end function


```

---

The idea behind our approach is to deactivate the connections  $i = 1 \dots, m_{l-1}$  that on average do not carry important information to the neuron  $j$  of layer  $l$  in comparison with other connections. As a result of Eq. 1, we simplify each neuron by using a smaller number



**Fig. 1** Neural network layers  $l - 1$  and  $l$  with  $m_{l-1} = 4$  and  $m_l = 2$  neurons, respectively. The weights associated to a connection between neurons  $i$  and  $j$  in layer  $l - 1$  to  $l$  are  $w_{ij}^{(l)}$



**Fig. 2** LeNet-300–100 architecture on MNIST before and after pruning, where connections are coloured with respect to importance score: blue (least important) → red (most important) (Color figure online)

of input connections, while causing minimal effective changes to the input signal of that neuron and subsequent minimal impact on the network.

Figure 2 shows the importance score of every connection in the LeNet-300–100 (LeCun et al., 1998) architecture before and after pruning when applied to the MNIST dataset. Connections are represented by thin blue lines when they have the lowest importance score  $s_{ij}^{(l)}$ , and go up to red thick lines when they are the most important. The figure demonstrates an interesting phenomenon besides the fact that there are much fewer connections after applying our pruning strategy: the importance scores of the connections become much closer to each other and there are no longer highly important and highly unimportant connections. We eliminate connections (contributors to the neuron’s signal) that on average do not contribute in terms of the strength of the signal that they bring to the neuron for a given dataset.

### 3.2 Convolutional layers

Our pruning approach for convolutional layers is similar to the one conducted on fully connected layers. We consider kernels and a bias in a particular filter as contributors to the signal produced by this filter.

#### Algorithm 2 Convolutional layers pruning

---

```

function CONV_PRUNING(network, X,  $\alpha$ )
2:    $\mathbf{X}^{(0)} \leftarrow \mathbf{X}$ 
   for conv layer  $l$  in CONV_Layers do
4:      $\mathbf{X}^{(l)} \leftarrow \text{conv\_layer}(\mathbf{X}^{(l-1)})$ 
     for every filter  $\mathbf{F}_j$  in conv_layer  $l$  do
6:       compute importance scores  $s_{ij} \forall$  kernel  $\mathbf{K}_{ij}^{(l)}$  and bias  $b_j^{(l)}$  in filter
        $\mathbf{F}_j^{(l)}$  using (2), (3).
        $\hat{s}_{ij}^{(l)} = \text{Sort}(s_{ij}^{(l)}, \text{order} = \text{descending})$ 
8:        $p_0 = \min\{p : \sum_{i=1}^p \hat{s}_{ij}^{(l)} \geq \alpha\}$ 
       prune kernel  $\mathbf{K}_{ij}^{(l)}$  with importance score  $s_{ij}^{(l)} < \hat{s}_{p_0j}^{(l)}$ 
10:    end for
   end for
12: return pruned network
end function

```

---

Assume we have  $m_{l-1}$ -channelled input samples  $\mathbf{X}^{(l-1)} = \{\mathbf{x}_1^{(l-1)}, \dots, \mathbf{x}_N^{(l-1)}\}$ , where  $\mathbf{x}_k^{(l-1)} = (x_{k1}^{(l-1)}, \dots, x_{km_{l-1}}^{(l-1)}) \in \mathbb{R}^{m_{l-1} \times h_{l-1}^1 \times h_{l-1}^2}$ , where  $h_{l-1}^1$  and  $h_{l-1}^2$  are the height and width of input images (or feature maps) for convolutional layer  $l$ . For every kernel  $\mathbf{K}_{1j}^{(l)}, \mathbf{K}_{2j}^{(l)}, \dots, \mathbf{K}_{m_jj}^{(l)}$ ,  $\mathbf{K}_{ij}^{(l)} = (k_{ijqt}^{(l)}) \in \mathbb{R}^{r_l \times r_l}$ ,  $1 \leq q, t \leq r_l$ ,  $r_l$  is a kernel size, and a bias  $b_j^{(l)}$  in filter  $\mathbf{F}_j^{(l)}$ , we define  $\hat{\mathbf{K}}_{ij}^{(l)} = \left( |k_{ijqt}^{(l)}| \right)$  as a matrix consisting of the absolute values of the matrix  $\mathbf{K}_{ij}^{(l)}$ .

Then we compute importance scores  $s_{ij}^{(l)}$ ,  $i \in \{1, 2, \dots, m_l\}$  of kernels  $\mathbf{K}_{ij}^{(l)}$  as follows:

$$s_{ij}^{(l)} = \frac{\frac{1}{N} \sum_{n=1}^N \left\| \hat{\mathbf{K}}_{ij}^{(l)} * |x_{ni}^{(l-1)}| \right\|_F}{S_j^{(l)}}, \quad (2)$$

$$s_{m_l+1,j}^{(l)} = \frac{|b_j^{(l)}| \sqrt{h_l^1 h_l^2}}{S_j^{(l)}}. \quad (3)$$

where  $S_j^{(l)} = \sum_{i=1}^{m_l} \left( \frac{1}{N} \sum_{n=1}^N \left\| \hat{\mathbf{K}}_{ij}^{(l)} * |x_{ni}^{(l-1)}| \right\|_F \right) + |b_j^{(l)}| \sqrt{h_l^1 h_l^2}$  is the total importance score in filter  $\mathbf{F}_j^{(l)}$  of layer  $l$ , and where  $*$  indicates a convolution operation, and  $\|\cdot\|_F$  the Frobenius norm.

In Eq. 2, we compute the amount of information that every kernel produces on average, analogously to what we do in fully connected layers.

## 4 Experiments

We test our pruning method for LeNet-300–100 and LeNet-5 (LeCun et al., 1998) on MNIST, VGG-19 (Simonyan and Zisserman, 2015) and VGG-like (Zagoruyko, 2015) on CIFAR-10/100 (Krizhevsky et al., 2009) and Tiny-ImageNet datasets. We evaluate our method on classification error, the percentage of pruned parameters  $\left( \frac{|w=0|}{|w|} (\%) \right)$  or the percentage of remaining parameters  $\left( \frac{|w \neq 0|}{|w|} (\%) \right)$  or compression rate  $\left( \frac{|w|}{|w \neq 0|} (\%) \right)$  and pruned FLOPs (floating-point operations), where  $|w \neq 0|$  refers to the number of unpruned connections. For the details about training and pruning hyperparameters, and FLOPs computation see Sections 1 and 2. We use the initialization method introduced in He et al. (2015) to initialize parameters. We run our experiments multiple times with different random initializations of the parameters. In our tables and figures, we present mean and standard deviation for the results averaged over these random initializations.

### 4.1 LeNets

We experiment with LeNets on the MNIST dataset. LeNet-300–100 is a fully connected network with 300 neurons and 100 neurons in two hidden layers respectively, and LeNet-5 Caffe, which is a modified version of LeCun et al. (1998), has two convolutional layers followed by one hidden layer and an output layer. We perform iterative pruning by retraining



**Table 1** Results for LeNet-300–100 and LeNet-5 trained and pruned on MNIST

Network	Method	Error (%)	Parameters retained (%)
LeNet-300–100	DNS (Guo et al., 2016)	1.99	1.79
	L-OBS (Dong et al., 2017)	1.96	1.5
	SWS (Ullrich et al., 2017)	1.94	4.3
	Sparse VD (Molchanov et al., 2017)	<b>1.92</b>	<b>1.47</b>
	<b>NNrelief (ours)</b>	1.98 ± 0.07	1.51 ± 0.07
LeNet-5	DNS (Guo et al., 2016)	0.91	0.93
	L-OBS (Dong et al., 2017)	1.66	0.9
	SWS (Ullrich et al., 2017)	0.97	0.5
	Sparse VD (Molchanov et al., 2017)	<b>0.75</b>	<b>0.36</b>
	<b>NNrelief (ours)</b>	0.97 ± 0.05	0.65 ± 0.02

Bold refers to the best results across all methods under consideration

For pruning, 1000 random training samples are chosen,  $\alpha_{fc} = 0.95$ ,  $\alpha_{conv} = 0.9$

the network starting from initial random parameters, instead of fine-tuning. Table 1 shows that our approach is among the best for both LeNets in terms of pruned parameters.

## 4.2 VGG

We perform our experiments on VGG-13 and VGG-like (adapted version of VGG-16 for CIFAR-10 dataset that has one fully connected layer less) on CIFAR-10, CIFAR-100 and Tiny-ImageNet datasets. We also show in this section the effect of two optimizers—Adam (Kingma and Ba, 2015) and SGD and two weight decay values on our pruning technique.

*CIFAR-10/100* Table 2 shows that after the reduction of the weights to less than 2% of the original number, the accuracy only drops by 0.1%, and that our approach outperforms others both in terms of pruned parameters and pruned FLOPs. In Fig. 3, we show the final architecture after pruning and the level of sparsity by layer. It can be observed that NNrelief with Adam optimizer achieves a high level of filter sparsity even though it prunes kernels.

We also perform pruning for standard VGG-13 on CIFAR-100 using 5 iterations, and considered three randomization seeds to provide reasonable statistical significance. The results from the different optimizers are presented in Fig. 4. The results show that Adam seems to perform the compression more aggressively, resulting in a higher compression rate, and also a slightly lower accuracy. The compression of about 14.5× after the fifth iteration without loss in accuracy is presented for Adam and weight decay  $5 \cdot 10^{-4}$ , while SGD allows to train a network with higher accuracy, but after 5 iterations we obtain smaller compression (about 4.5×) and the decrease of accuracy by 0.5% with the same weight decay. We obtain a lower compression rate for both optimizers when the weight decay is lower ( $10^{-4}$ ). Figure 5 shows the final architecture and sparsity that we achieve with Adam optimizer at the end of the pruning procedure. We observe a high level of filter sparsity with Adam, as also reported in Mehta et al. (2019), but good compression is also obtained for the SGD optimizer.

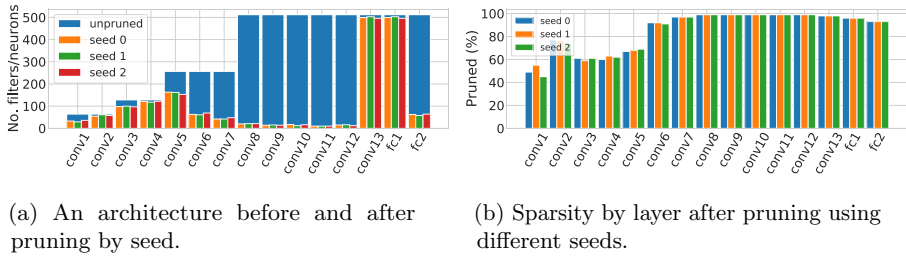
**Tiny-ImageNet** We also apply our approach to Tiny-ImageNet dataset that is a subset of the ImageNet dataset with 200 classes and an image spatial resolution of 64×64. Following SNIP strategy, we use strides [2, 2] in the first convolutional layer to reduce the size of

**Table 2** Results for VGG-like trained on CIFAR-10 with Adam optimizer.

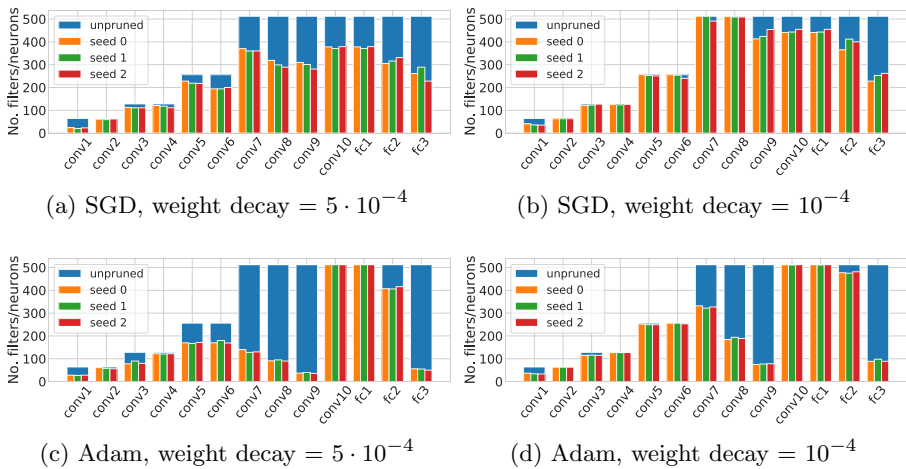
Network	Method	Acc (%) (baseline)	Acc drop (%)	Parameters retained (%)	FLOPs pruned (%)
VGG-like	Pruning (Li et al., 2017)	93.25	-0.15	36	34.2
	Sparse VD (Molchanov et al., 2017)	92.3	0.0	2.1	N/A
	BC-GNJ (Louizos et al., 2017)	91.6	0.2	6.7	55.6
	BC-GHS (Louizos et al., 2017)	91.6	0.6	5.5	61.7
	SNIP (Lee et al., 2019)	91.7	-0.3	3.0	N/A
	HRank (Lin et al., 2020)	93.96	1.62	17.9	65.3
	CHIP (Sui et al., 2021)	93.96	0.24	16.7	66.6
	<b>NNrelief (ours)</b>	<b>92.5</b>	<b>0.1</b>	<b>1.92 ± 0.02</b>	<b>75.5 ± 0.4</b>

**Bold** refers to the best results across all methods under consideration

During retraining 60 epochs are used; the learning rate is decreased by 10 every 20 epochs. For NNrelief (ours)  $\alpha_{\text{conv}} = \alpha_{\text{fc}} = 0.95$  and 1000 samples are chosen for the importance scores computation. Sign '±' means that accuracy has increased after pruning



**Fig. 3** Architecture structure for VGG-like on CIFAR-10 with Adam optimizer considering three random initializations



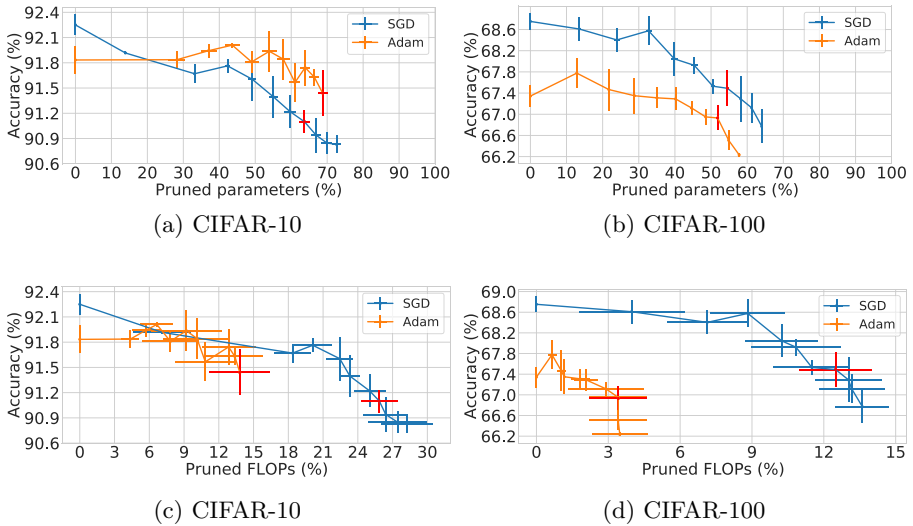
**Fig. 4** Results for VGG-13 over three seeds; mean values are used to compute dots and standard deviation are shown with error bars. We compare two optimizers, SGD and Adam, and two different values of weight decay for evaluation after 5 pruning iterations

images. The results are reported in Table 3. We observe a compression by more than 40 times almost without any loss of accuracy ( $-0.03\%$ ) using Adam optimizer.

### 4.3 ResNets

In addition, we test our approach on ResNet architectures for CIFAR-10/100 and Tiny-ImageNet datasets. Our main objective is to prune as many parameters as we can without significant loss of accuracy. For ResNet-20/56 on CIFAR-10, we perform iterative pruning over 10 iterations, training the model with SGD and Adam. In Table 4 we compare our results with other approaches, and Fig. 6 displays the pruning history. We achieve a higher percentage of pruned parameters for ResNet-20 and ResNet-56, using both optimizers with comparable final accuracy.

A notable observation for some seeds when training with SGD is that separate convolutional blocks are pruned and only residual blocks remain, i.e. the signal propagates through



**Fig. 5** VGG-13 architecture on CIFAR-100 trained with SGD (top) and Adam (bottom), and weight decay equal to  $5 \cdot 10^{-4}$  (left) and  $10^{-4}$  (right) and pruned with 5 iterations. The results for three different seeds are presented

**Table 3** Results for VGG-like trained on Tiny-ImageNet with Adam optimizer.

Network	Method	Acc (%) (baseline)	Acc drop (%)	Parameters retained (%)	FLOPs pruned (%)
VGG-like	SNIP (Lee et al., 2019)	45.14	0.87	5.0	N/A
	<b>NNrelief (ours)</b>	45.63	0.03	<b>2.32 ± 0.06</b>	<b>75 ± 0.32</b>

Bold refers to the best results across all methods under consideration

During retraining 50 epochs are used; the learning rate is decreased by 10 every 20 epochs. For NNrelief (ours)  $\alpha_{\text{conv}} = \alpha_{\text{fc}} = 0.95$  and 2000 samples are chosen for the importance scores computation

skip connections and not through main convolutional layers since they do not contribute to the sum (Fig. 8).

For ResNet-20/56 on CIFAR-100 the results are presented in Table 5 and by iteration in Fig. 7.

**Tiny-ImageNet** Training of a modified ResNet-18 with 16, 32, 64 and 128 output channels indicates that we can prune more than 50% of the parameters with both optimizers (see Fig. 9). Adam, however, maintains a higher level of accuracy during retraining than SGD. We use 2000 samples to evaluate importance scores, and weight decay for (re-)training equal to  $5 \cdot 10^{-4}$ .

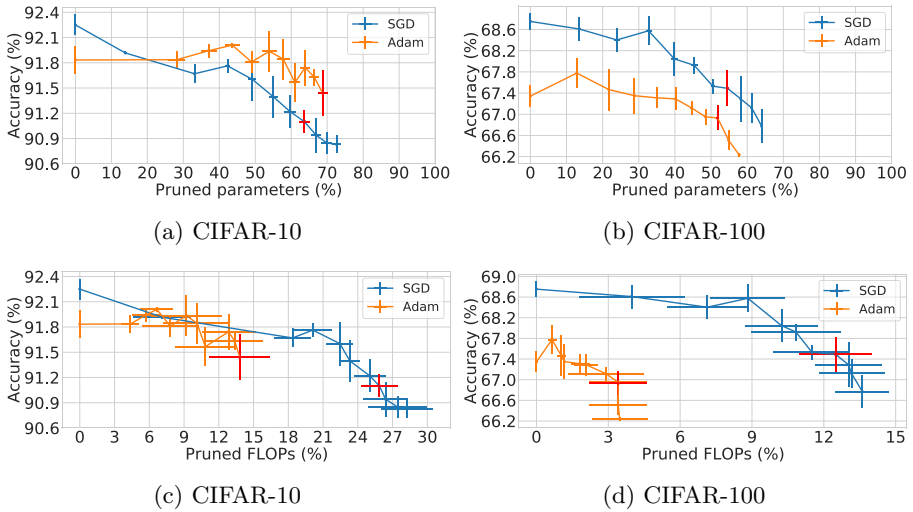
**Sparsity** In contrast to VGG results, we do not observe a significant difference between optimizers in terms of produced sparsity for CIFAR-10/100 and Tiny-ImageNet. In the works that explore the question of sparsity (Mehta et al., 2019), ResNet architecture was not explored. It seems that the skip connections allow the signal to pass unhindered to the next layers, making it harder for the Adam optimizer to find which activations are significant, and aggressively prune the less significant connections.

**Table 4** Results for ResNets trained on CIFAR-10 with SGD and Adam optimizers.

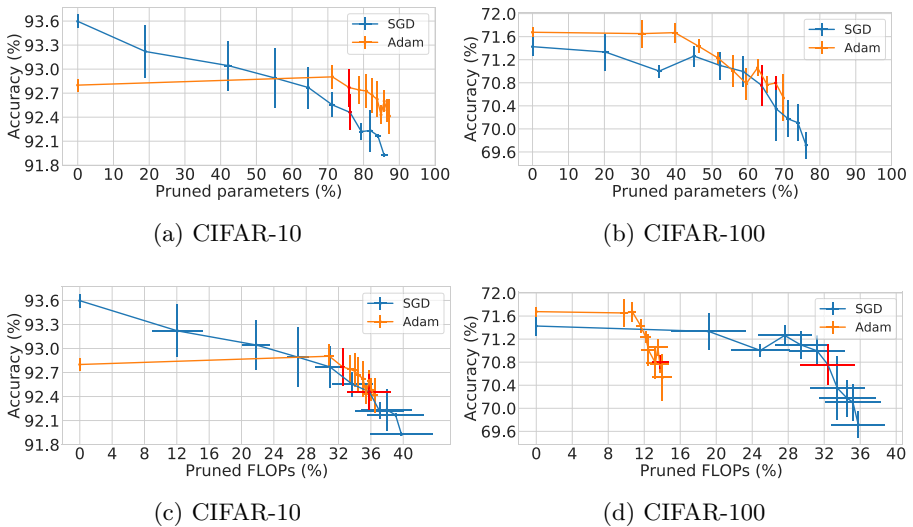
ResNet	Method	Acc (%) (baseline)	Acc drop (%)	Parameters pruned (%)	FLOPs pruned (%)
20	SFP (He et al., 2018)	92.2	1.37	30	42.2
	SNLI (Ye et al., 2018)	92.0	1.1	37.2	N/A
	SSS (Huang and Wang, 2018)	92.8	2	45	<b>60</b>
	CNN-CFC (Li et al., 2019)	92.2	1.07	42.75	41.6
	<b>NNrelief</b> (SGD)	92.25 ± 0.12	1.15 ± 0.13	<b>63.68 ± 1.52</b>	25.85 ± 1.57
	<b>NNrelief</b> (Adam)	91.83 ± 0.16	<b>0.39 ± 0.27</b>	<b>68.75 ± 1.26</b>	13.81 ± 2.6
	Pruning-B (Li et al., 2017)	93.04	-0.02	13.7	27.6
	SFP (He et al., 2018)	93.59	-0.19	30	41.1
	NISP (Yu et al., 2018)	N/A	0.03	42.6	43.61
	CNN-CFC (Li et al., 2019)	93.14	-0.24	43.09	42.78
56	HRank (Lin et al., 2020)	93.14	1.22	69.74	70.9
	CHIP (Sui et al., 2021)	93.36	0.09	42.4	50
	<b>NNrelief</b> (SGD)	93.36	2.54	68.1	<b>74.1</b>
	<b>NNrelief</b> (Adam)	93.26	1.21	71.8	72.3
	Pruning-B (Li et al., 2017)	93.6 ± 0.08	1.14 ± 0.22	<b>76.2 ± 0.44</b>	35.8 ± 2.69
	<b>NNrelief</b> (SGD)	92.8 ± 0.08	0.03 ± 0.23	<b>75.95 ± 0.22</b>	32.5 ± 0.31
	HRank (Lin et al., 2020)	93.14	1.22	69.74	70.9
	CHIP (Sui et al., 2021)	93.36	0.09	42.4	50
	<b>NNrelief</b> (SGD)	93.36	2.54	68.1	<b>74.1</b>
	<b>NNrelief</b> (Adam)	93.26	1.21	71.8	72.3

**Bold** refers to the best results across all methods under consideration

During retraining 60 epochs are used; the learning rate is decreased by 10 every 20 epochs. For NNrelief (ours)  $\alpha_{\text{conv}} = 0.95, \alpha_{\text{fc}} = 0.99$




**Fig. 6** ResNet-20 on CIFAR-10 and CIFAR-100 by iteration averaged over three random initializations, the error bars represent standard deviation for these three runs. The red dot corresponds to the selected best iteration



**Fig. 7** ResNet-56 on CIFAR-10 and CIFAR-100 by iteration averaged over three seeds, the error bars represent standard deviation for these three runs. The red dot corresponds to the selected best iteration

## 5 Discussion

To illustrate the meaning of the importance scores we compare our approach with the magnitude-based one. Also, we present error bounds for the difference of a signal in the layer before and after pruning.

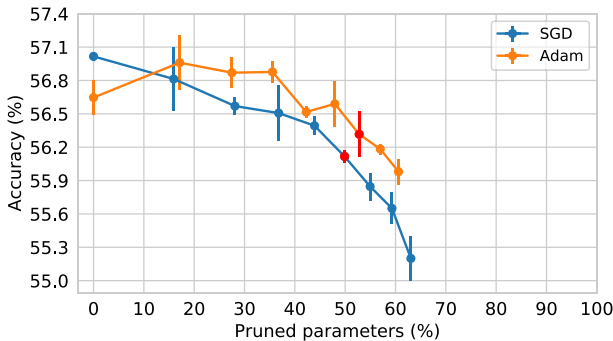


**Fig. 8** ResNet-56 architecture on CIFAR-10 by seed with SGD optimizer

**Table 5** Results for ResNets trained on CIFAR-100 with SGD and Adam optimizers

ResNet	Optimizer	Acc (%) (baseline)	Acc drop (%)	Parameters pruned (%)	FLOPs pruned (%)
20	SGD	68.75 ± 0.15	1.26 ± 0.33	54.54 ± 0.24	12.49 ± 1.52
	Adam	67.34 ± 0.2	0.41 ± 0.23	51.82 ± 0.23	3.42 ± 1.21
56	SGD	71.43 ± 0.16	0.43 ± 0.26	58.62 ± 0.26	31.2 ± 3.09
	Adam	71.68 ± 0.09	0.88 ± 0.12	67.72 ± 0.29	13.79 ± 0.94

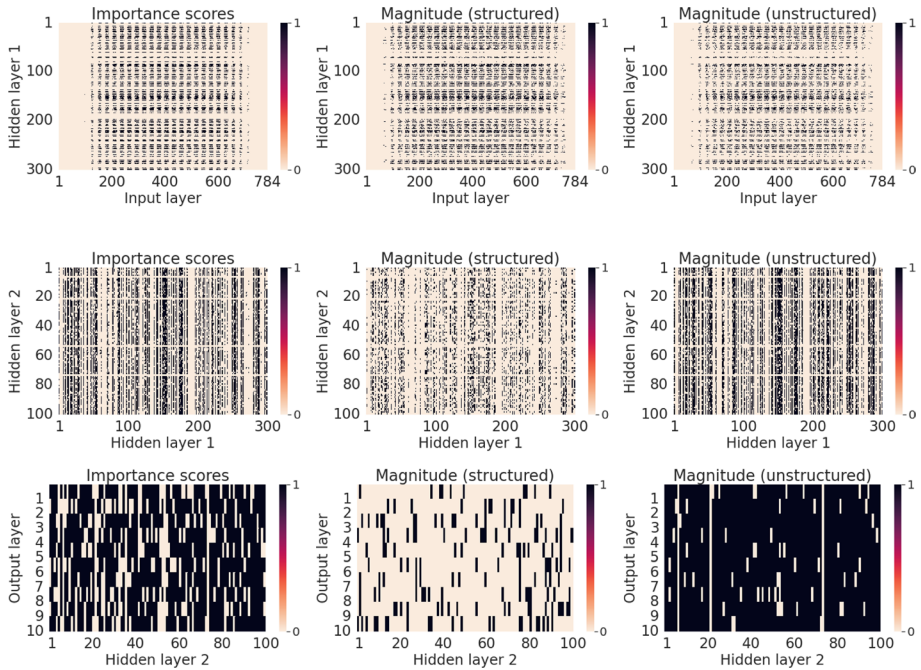
During retraining 80 epochs are used; the learning rate is decreased by 10 every 30 epochs. For NNrelief (ours)  $\alpha_{conv} = 0.95, \alpha_{fc} = 0.99$



**Fig. 9** ResNet-18 on Tiny-ImageNet trained with SGD and Adam optimizers. The results are averaged over three random initializations. The bars represent standard deviations considering three runs

### 5.1 Comparison with magnitude-based approach

In order to show the difference between our approach and the magnitude-based approach (Han et al., 2015), we build the heatmaps with the location of top 15.5% of the most significant connections in the LeNet-300–100 pretrained on MINIST for both our proposed method and the magnitude-based pruning, since this number corresponds to the percentage of parameters that remain after pruning with our proposed rule (importance scores) using  $\alpha = 0.95$ . The importance scores (IS) are computed from Eq. 1, while for the magnitude-based rule, we consider both structured and unstructured pruning. In the structured case, we select the top 15.5% in each layer independently, and in the unstructured, the top 15.5% from all layers are selected, meaning that some layers may contain more or less than 15.5% of the parameters.



**Fig. 10** The remaining 15.5% connections according to Importance scores and magnitude-based rule (structured and unstructured pruning) for trained LeNet-300–100 on MNIST

**Table 6** Jaccard index (similarity) between IS rule ( $\alpha = 0.95$ ) and magnitude-based one (structured and unstructured) for each layer if 84.5% of parameters are pruned in LeNet-300–100 pretrained on MNIST

	Layer	Magnitudes (structured)	Magnitudes (unstructured)
Importance scores	784 → 300	0.08	0.08
	300 → 100	0.11	0.18
	100 → 10	0.16	0.66

**Table 7** Number of active neurons for magnitude-based pruning and with importance scores one (NNrelief) for LeNet-300–100 pretrained on MNIST

Method	Architecture
Original	784 → 300 → 100 → 10
Importance scores	380 → 131 → 96 → 10
Magnitude-based (unstructured)	530 → 136 → 97 → 10

Figure 10 displays the difference between our approach and the magnitude-based (structured and unstructured) one when determining the significance of network connections. According to the figure, we can see different patterns for the remaining connection, especially for the unstructured magnitude-based rule, where more connections are retained in the last layer compared to the IS rule. Overall, we present the Jaccard index (Jaccard, 1912) between IS and both magnitude-based rules in Table 6. The Jaccard index (or Intersection over Union, IoU) for two sets  $A$  and  $B$  is defined as follows:  $\text{IoU} = \frac{|A \cap B|}{|A \cup B|}$ .



Both Fig. 10 and Table 6 demonstrate that the IS rule prunes the network differently from both magnitude-based rules. First, the binary heatmap patterns in the figure indicate a clear difference between the methods. In addition, the IoU indicators are low for every case, meaning that the IS pruning proposed herein determines the most significant connections differently to both magnitude-based rules. From empirical results, we observe that IS-based pruning deactivates more connections in overparameterized layers. In Table 7, we present the number of active neurons by layer after one pruning iteration (15.5% of all parameters are retained). As it can be seen, the IS metric deactivates more neurons in every layer compared to unstructured magnitude pruning (Han et al., 2015; Frankle and Carbin, 2019). This happens because some neurons produce a low signal (or zero signal because of the ReLU activation function).

In contrast to magnitude-based pruning, NNrelief automatically determines the number of connections to prune when  $\alpha$  is given. However, for magnitude-based pruning, the number of pruned connections is a hyperparameter which is difficult to set. For example, if we prune 80% of the remaining parameters at every iteration as discussed in the lottery ticket hypothesis work (Frankle and Carbin, 2019), then after 11 iterations magnitude-based pruning retains  $100\% \cdot 0.8^{11} \approx 8.6\%$  of the parameters, while NNrelief retains 1.51% (see Tables 1 and 10). This is achieved due to the usage of the input signal to estimate the importance scores, which helps to determine the number of connections to prune.

### 5.2 Error estimation

To complement the analysis of our approach we derive error bounds by considering the difference between a signal in the trained neuron and a reduced one without retraining, before and after the activation function. This analysis is important to make sure NNrelief does not destroy the pretrained structure, and that changes in the output neurons are not dramatic enough to prevent successful retraining. To illustrate this, we compare the derived error bounds with the observed changes in ResNet-20 pretrained on CIFAR-10 if the fully connected layer is pruned. Without loss of generality we may assume the first  $p$  connections in the neuron  $j$  of layer  $l$  are kept and that the bias  $b_j^{(l)}$  as well as weights  $w_{p+1,j}^{(l)}, \dots, w_{m_j,l}^{(l)}$  are pruned in a particular neuron  $j$ . Then, for any input  $x_n^{(l-1)} \in \mathbf{X}^{(l-1)} = \{x_1^{(l-1)}, \dots, x_N^{(l-1)}\} \subset \mathbb{R}^{m_{l-1}}$  we obtain:

$$\begin{aligned} \delta_p(x_n^{(l-1)}) &= \left| \sum_{k=1}^{m_{l-1}} w_{kj}^{(l)} x_{nk}^{(l-1)} + b_j^{(l)} - \sum_{k=1}^p w_{kj}^{(l)} x_{nk}^{(l-1)} \right| = \left| \sum_{k=p+1}^{m_{l-1}} w_{kj}^{(l)} x_{nk}^{(l-1)} + b_j^{(l)} \right| \\ &\leq |b_j^{(l)}| + \sum_{k=p+1}^{m_{l-1}} \left| w_{kj}^{(l)} x_{nk}^{(l-1)} \right|. \end{aligned} \tag{4}$$

By averaging over all samples:

$$\begin{aligned} \overline{\delta}_p &\leq |b_j^{(l)}| + \sum_{k=p+1}^{m_{l-1}} \overline{|w_{kj}^{(l)} x_k^{(l-1)}|} = S_j^{(l)} \cdot \left( \frac{|b_j^{(l)}|}{S_j^{(l)}} + \frac{\sum_{k=p+1}^{m_{l-1}} \overline{|w_{kj}^{(l)} x_k^{(l-1)}|}}{S_j^{(l)}} \right) \\ &= S_j^{(l)} \cdot \sum_{k=p+1}^{m_{l-1}} s_{kj}^{(l)} = S_j^{(l)} (1 - \alpha). \end{aligned} \tag{5}$$

Therefore, from (5) we see that the sequence of residuals  $\{\bar{\delta}_p\}_{p=1}^m$  decreases with respect to  $p$  and increases with  $\alpha$ . Also, for any Lipschitz continuous activation function  $\varphi \in \text{Lip}(X)$ ,  $X \subset \mathbb{R}$  with Lipschitz constant  $C$  we obtain:

$$\begin{aligned} \Delta_p(x_n^{(l-1)}) &= \left| \varphi \left( \sum_{k=1}^{m_{l-1}} w_{kj}^{(l)} x_{nk}^{(l-1)} + b_j^{(l)} \right) - \varphi \left( \sum_{k=1}^p w_{kj}^{(l)} x_{nk}^{(l-1)} \right) \right| \\ &\leq C \left| \sum_{k=p+1}^{m_{l-1}} w_{kj}^{(l)} x_{nk}^{(l-1)} + b_j^{(l)} \right| = C \delta_p(x_n^{(l-1)}). \end{aligned} \quad (6)$$

As a result, from Eqs. (5) and (6):

$$\bar{\Delta}_p \leq C S_j^{(l)} (1 - \alpha). \quad (7)$$

For the point-wise estimation on set  $\mathbf{X}^{(l-1)} = \{x_1^{(l-1)}, \dots, x_N^{(l-1)}\} \subset \mathbb{R}^{m_{l-1}}$ ,  $\forall x_i^{(l-1)}$  we obtain:

$$\Delta_p(x_i^{(l-1)}) \leq \max_{x_n^{(l-1)} \in \mathbf{X}^{(l-1)}} C \delta_p(x_n^{(l-1)}). \quad (8)$$

Since ReLU, ELU, sigmoid and, tanh are Lipschitz continuous functions then (7) and (8) hold for them, meaning that the sequence of residuals  $\{\bar{\Delta}_p\}_{p=1}^{m_l}$  decreases as well with the increase of  $\alpha$ . Moreover, in the case of ReLU and ELU, the Lipschitz constant is  $C = 1$ . In summary, by increasing  $\alpha$  we indeed reduce the approximation error of the neuron.

Analogously, for input  $m_{l-1}$ -channeled pruning set  $\mathbf{X}^{(l-1)} = \{\mathbf{x}_1^{(l-1)}, \dots, \mathbf{x}_N^{(l-1)}\}$ , where  $\mathbf{x}_k^{(l-1)} = (x_{k1}^{(l-1)}, \dots, x_{km}^{(l-1)}) \in \mathbb{R}^{m_{l-1} \times h_{l-1}^1 \times h_{l-1}^2}$ , where  $h_{l-1}^1$  and  $h_{l-1}^2$  are height and width of input images (or feature maps) for convolutional layer  $l-1$ , and  $\mathbf{B}_j^{(l)} \in \mathbb{R}^{h_l^1 \times h_l^2}$  is a matrix that consists of the same bias value  $b_j^{(l)}$  at every element we obtain:

$$\begin{aligned} \delta_p(x_k^{(l-1)}) &= \left\| \sum_{i=1}^{m_{l-1}} \mathbf{K}_i^{(l)} * x_{ki}^{(l-1)} + \mathbf{B}_j^{(l)} - \sum_{i=1}^p \mathbf{K}_i^{(l)} * x_{ki}^{(l-1)} \right\|_F \\ &= \left\| \sum_{i=p+1}^{m_{l-1}} \mathbf{K}_i^{(l)} * x_{ki}^{(l-1)} + \mathbf{B}_j^{(l)} \right\|_F \leq \sum_{i=p+1}^{m_{l-1}} \left\| \mathbf{K}_i^{(l)} * x_{ki}^{(l-1)} \right\|_F + \left\| \mathbf{B}_j^{(l)} \right\|_F \\ &\leq \sum_{i=p+1}^{m_{l-1}} \left\| \mathbf{K}_i^{(l)} * |x_{ki}^{(l-1)}| \right\|_F + \sqrt{h_l^1 h_l^2} |b_j^{(l)}|. \end{aligned} \quad (9)$$

Then, by averaging over all pruning samples similarly to the fully connected case we obtain:

$$\bar{\delta}_p \leq \sum_{i=p+1}^{m_{l-1}} \frac{1}{N} \sum_{k=1}^N \left\| \mathbf{K}_i^{(l)} * x_{ki}^{(l-1)} \right\|_F = S_j^{(l)} (1 - \alpha), \quad (10)$$

from which for any Lipschitz continuous activation functions  $\varphi$  with Lipschitz constant  $C$ , we obtain:

$$\begin{aligned} \bar{\Delta}_p &= \frac{1}{N} \sum_{k=1}^N \left\| \varphi \left( \sum_{i=1}^{m_{l-1}} \mathbf{K}_i^{(l)} * x_{ki}^{(l-1)} + \mathbf{B}_j^{(l)} \right) - \varphi \left( \sum_{i=1}^p \mathbf{K}_i^{(l)} * x_{ki}^{(l-1)} \right) \right\|_F \\ &\leq C S_j^{(l)} (1 - \alpha). \end{aligned} \tag{11}$$

Inequalities (7), (11) show how the signal in a neuron or filter, respectively, changes after the layer is pruned. Let us consider how this change influences the network output. We define  $N_{i_l}^{(l)}(x; \mathbf{W}^{(1:l)})$  as a value in the  $i_l^{\text{th}}$  neuron in layer  $l$ , where  $\mathbf{W}^{(1:l)} = (\mathbf{W}^{(1)}, \dots, \mathbf{W}^{(l)})$  are parameters in layers  $1, \dots, l$ . The matrix that we obtain after pruning layer  $l$  we define as  $\hat{\mathbf{W}}^{(l)}$ , and  $m_l$  is the number of neurons in layer  $l$ ,  $1 \leq l \leq L$ . Then, for  $1 \leq l < L$  (for  $l = L$  see (5) in Section 5.2):

$$\begin{aligned} \varepsilon(x) &= \left| N_{i_L}^{(L)}(x; \mathbf{W}^{(1:L)}) - N_{i_L}^{(L)}(x; \mathbf{W}^{(1:l-1)}, \hat{\mathbf{W}}^{(l)}, \mathbf{W}^{(l+1:L)}) \right| \\ &= \left| \sum_{i_{L-1}=1}^{m_{L-1}} w_{i_{L-1}i_L}^{(L)} \left[ \varphi \left( N_{i_{L-1}}^{(L-1)}(x; \mathbf{W}^{(1:L-1)}) \right) - \varphi \left( N_{i_{L-1}}^{(L-1)}(x; \mathbf{W}^{(1:l-1)}, \hat{\mathbf{W}}^{(l)}, \mathbf{W}^{(l+1:L-1)}) \right) \right] \right| \\ &\leq C \sum_{i_{L-1}=1}^{m_{L-1}} \left| w_{i_{L-1}i_L}^{(L)} \left| N_{i_{L-1}}^{(L-1)}(x; \mathbf{W}^{(1:L-1)}) - N_{i_{L-1}}^{(L-1)}(x; \mathbf{W}^{(1:l-1)}, \hat{\mathbf{W}}^{(l)}, \mathbf{W}^{(l+1:L-1)}) \right| \right| \end{aligned} \tag{12}$$


Performing similar transformations as in (12) until layer  $l$ , we obtain:

$$\begin{aligned} \varepsilon(x) &\leq C^{L-l} \sum_{i_{L-1}=1}^{m_{L-1}} |w_{i_{L-1}i_L}^{(L)}| \sum_{i_{L-2}=1}^{m_{L-2}} |w_{i_{L-2}i_{L-1}}^{(L-1)}| \dots \sum_{i_j=1}^{m_j} |w_{i_j i_{j+1}}^{(l)}| |\varphi(N_{i_j}^{(l)}(x; \mathbf{W}^{(1:l)})) \\ &\quad - \varphi(N_{i_j}^{(l)}(x; \mathbf{W}^{(1:l-1)}, \hat{\mathbf{W}}^{(l)})| \end{aligned} \tag{13}$$

Then, averaging (13) over all pruning samples from  $\mathbf{X}$  and using (7):

$$\begin{aligned} \bar{\varepsilon} &\leq C^{L-l} \sum_{i_{L-1}=1}^{m_{L-1}} |w_{i_{L-1}i_L}^{(L)}| \sum_{i_{L-2}=1}^{m_{L-2}} |w_{i_{L-2}i_{L-1}}^{(L-1)}| \dots \sum_{i_j=1}^{m_j} |w_{i_j i_{j+1}}^{(l)}| |\bar{\Delta}_{p_{i_j}}^{(l)}| \\ &\leq (1 - \alpha) C^{L-l} \sum_{i_{L-1}=1}^{m_{L-1}} |w_{i_{L-1}i_L}^{(L)}| \sum_{i_{L-2}=1}^{m_{L-2}} |w_{i_{L-2}i_{L-1}}^{(L-1)}| \dots \sum_{i_j=1}^{m_{j+1}} |w_{i_j i_{j+1}}^{(l)}| |S_{i_j}^{(l)}|. \end{aligned} \tag{14}$$

Inequality (14) shows the bounds of the network error change after pruning on the pruning set  $\mathbf{X}$ . Fig. 11 shows the absolute changes in output responses for ResNet-20 after pruning on CIFAR-10 with  $\alpha_{fc} = 0.95$ . The error bars (in black) show the theoretical bounds of the difference estimated in (8) and (13), and the dots demonstrate the actual changes computed on the test data. From the figure, one can observe that the theoretical bounds capture the signal changes well in most cases, especially, for output neurons 0, 6 and 7.



**Fig. 11** Absolute difference in the neurons' response at the output layer before and after pruning ResNet-20 ( $\alpha_{fc} = 0.95$ ). The bars represent the theoretical bounds, and dots present empirical observations for 10,000 images

## 6 Conclusion

We propose a new pruning strategy and associated error estimation that is applicable to different neural network architectures. Our algorithm is based on eliminating layers' kernels or connections that do not contribute to the signal in the next layer. In this method, we estimate the average signal strength coming through each connection via proposed importance scores and try to keep the signal level in the neurons close to the original level. Thus, the number of removed parameters in a layer is selected adaptively, rather than being predetermined. As a result, the algorithm obtains a subnetwork capable of propagating most of the original signal while using fewer parameters than other reported strategies. We explore the effect of Adam and SGD optimizers on our pruning strategy. The results show that for VGG architecture Adam optimizer gives much higher compression than SGD, however, for ResNet, we do not observe similar behaviour. We observe that the obtained subnetwork tends to homogenize connection importance, hinting that every remaining connection is approaching similar importance for the dataset on average.

## Appendix A: Pruning setups

We implement our approach with PyTorch (Paszke et al., 2019). Table 8 shows the training parameters for each network.

Table 9 shows the parameters for retraining on every iteration.

Table 10 shows the pruning parameters and total number of iterations.

**Table 8** Training setups

Network	Dataset	Optimizer	Learning rate (by epoch)	Weight decay
LeNet-5/300–100	MNIST	Adam	$\begin{cases} 10^{-3} & 1 \leq epoch \leq 30 \\ 10^{-4} & 31 \leq epoch \leq 60 \end{cases}$	$5 \cdot 10^{-4}$
VGG-like/13	CIFAR-10/100, Tiny-ImageNet	SGD/Adam	$\begin{cases} 10^{-1}/10^{-3} & 1 \leq epoch \leq 80 \\ 10^{-2}/10^{-4} & 81 \leq epoch \leq 120 \\ 10^{-3}/10^{-5} & 121 \leq epoch \leq 150 \end{cases}$	$5 \cdot 10^{-4}$
ResNet-20/56	CIFAR-10/100	SGD	$\begin{cases} 0.1 & 1 \leq epoch \leq 80 \\ 0.01 & 81 \leq epoch \leq 120 \\ 0.001 & 121 \leq epoch \leq 150 \end{cases}$	$5 \cdot 10^{-4}$
		Adam	$\begin{cases} 10^{-3} & 1 \leq epoch \leq 120 \\ 10^{-4} & 121 \leq epoch \leq 160 \\ 10^{-5} & 161 \leq epoch \leq 200 \end{cases}$	$5 \cdot 10^{-4}$

**Table 9** Retraining parameters

Network	Dataset	Optimizer	Learning rate (by epoch)	Weight decay
LeNet-5/300–100	MNIST	Adam	$\begin{cases} 10^{-3} & 1 \leq epoch \leq 30 \\ 10^{-4} & 31 \leq epoch \leq 60 \end{cases}$	$5 \cdot 10^{-4}$
VGG-like/13	CIFAR-10/100, Tiny-ImageNet	SGD/Adam	$\begin{cases} 10^{-1}/10^{-3} & 1 \leq epoch \leq 20 \\ 10^{-2}/10^{-4} & 21 \leq epoch \leq 40 \\ 10^{-3}/10^{-5} & 41 \leq epoch \leq 60 \end{cases}$	$5 \cdot 10^{-4}$
ResNet-20/56	CIFAR-10	SGD/Adam	$\begin{cases} 10^{-1}/10^{-3} & 1 \leq epoch \leq 20 \\ 10^{-2}/10^{-4} & 21 \leq epoch \leq 40 \\ 10^{-3}/10^{-5} & 41 \leq epoch \leq 60 \end{cases}$	$5 \cdot 10^{-4}$
	CIFAR-100	SGD/Adam	$\begin{cases} 10^{-1}/10^{-3} & 1 \leq epoch \leq 30 \\ 10^{-2}/10^{-4} & 21 \leq epoch \leq 60 \\ 10^{-3}/10^{-5} & 41 \leq epoch \leq 80 \end{cases}$	$5 \cdot 10^{-4}$

**Table 10** Pruning parameters

Network	Dataset	$(\alpha_{\text{conv}}, \alpha_{\text{fc}})$	# iterations	Best iteration SGD/Adam
LeNet-300–100	MNIST	(-, 0.95)	15	11
LeNet-5	MNIST	(0.9,0.95)	20	20
VGG-like	CIFAR-10	(0.95,0.95)	6	6
	Tiny-ImageNet	(0.95,0.95)	5	5
ResNet-20 (SGD/Adam)	CIFAR-10	(0.95,0.99)	10	7/10
	CIFAR-100	(0.95,0.99)	10	7/8
ResNet-56 (SGD/Adam)	CIFAR-10	(0.95,0.99)	10	6/2
	CIFAR-100	(0.95,0.99)	10	6/9

## Appendix B: FLOPs computation

According to Molchanov et al. (2016), we compute FLOPs as follows:

- for a convolutional layer:  $\text{FLOPs} = 2HW(C_{in}K^2 + 1)C_{out}$  where  $H$ ,  $W$  and  $C_{in}$  are height, width and number of channels of the input feature map,  $K$  is the kernel width (and height due to symmetry), and  $C_{out}$  is the number of output channels.
- for a fully connected layer:  $\text{FLOPs} = (2I - 1)O$ , where  $I$  is the input dimensionality and  $O$  is the output dimensionality.

**Author contributions** Not applicable

**Funding** Not applicable

**Data availability** Not applicable

**Code availability** PyTorch [1] implementation of the code is available at: <https://github.com/adekhovich/NNrelief>.

## Declarations

**Conflict of interest** The authors have no conflicts of interest to declare.

**Ethical approval** Not applicable

**Consent to participate** Not applicable

**Consent for publication** Not applicable

## References

- Aghasi, A., Abdi, A., Nguyen, N., & Romberg, J. (2017). Net-trim: Convex pruning of deep neural networks with performance guarantee. *Advances in Neural Information Processing Systems*, 30, 3177–3186.

- Ahmad, S., & Scheinkman, L. (2019). How can we be so dense? The benefits of using highly sparse representations. arXiv preprint [arXiv:1903.11257](https://arxiv.org/abs/1903.11257)
- Ancona, M., Öztireli, C., & Gross, M. (2020). Shapley value as principled metric for structured network pruning. arXiv preprint [arXiv:2006.01795](https://arxiv.org/abs/2006.01795)
- Beck, A., & Teboulle, M. (2009). A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM Journal on Imaging Sciences*, 2(1), 183–202.
- Dong, X., Chen, S., & Pan, S. (2017). Learning to prune deep neural networks via layer-wise optimal brain surgeon. In *Advances in neural information processing systems* (pp. 4857–4867).
- Frankle, J., & Carbin, M. (2019). The lottery ticket hypothesis: Finding sparse, trainable neural networks. In *International conference on learning representations (ICLR)*
- Garg, Y., & Candan, K.S. (2020). isparse: Output informed sparsification of neural network. In *Proceedings of the 2020 international conference on multimedia retrieval* (pp. 180–188).
- Geng, L., & Niu, B. (2022). Pruning convolutional neural networks via filter similarity analysis. *Machine Learning*, 111(9), 3161–3180.
- Guo, Y., Yao, A. & Chen, Y. (2016). Dynamic network surgery for efficient DNNs. In *Advances in neural information processing systems* (pp. 1379–1387).
- Han, S., Mao, H., & Dally, W.J. (2016). Deep compression: Compressing deep neural networks with pruning, trained quantization and Huffman coding. In *International conference on learning representations (ICLR)*
- Han, S., Pool, J., Tran, J., & Dally, W. (2015). Learning both weights and connections for efficient neural network. In *Advances in neural information processing systems* (pp. 1135–1143).
- Hassibi, B., & Stork, D.G. (1993). Second order derivatives for network pruning: Optimal brain surgeon. In *Advances in neural information processing systems* (pp. 164–171)
- He, Y., Zhang, X., & Sun, J. (2017). Channel pruning for accelerating very deep neural networks. In *Proceedings of the IEEE international conference on computer vision* (pp. 1389–1397).
- He, Y., Kang, G., Dong, X., Fu, Y., & Yang, Y. (2018). Soft filter pruning for accelerating deep convolutional neural networks. In *International joint conference on artificial intelligence (IJCAI)* (pp. 2234–2240).
- He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 770–778).
- He, K., Zhang, X., Ren, S., & Sun, J. (2015). Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In: *Proceedings of the IEEE international conference on computer vision* (pp. 1026–1034).
- Hu, H., Peng, R., Tai, Y.-W., & Tang, C.-K. (2016). Network trimming: A data-driven neuron pruning approach towards efficient deep architectures. arXiv preprint [arXiv:1607.03250](https://arxiv.org/abs/1607.03250)
- Huang, Z., & Wang, N. (2018). Data-driven sparse structure selection for deep neural networks. In *Proceedings of the European conference on computer vision (ECCV)* (pp. 304–320).
- Jaccard, P. (1912). The distribution of the flora in the alpine zone. 1. *New Phytologist*, 11(2), 37–50.
- Karnin, E. D. (1990). A simple procedure for pruning back-propagation trained neural networks. *IEEE Transactions on Neural Networks*, 1(2), 239–242.
- Kingma, D.P. & Ba, J. (2015). Adam: A method for stochastic optimization. In *3rd international conference on learning representations, ICLR*.
- Krizhevsky, A. (2009). Learning multiple layers of features from tiny images. Master's thesis, Department of Computer Science, University of Toronto.
- LeCun, Y., Bottou, L., Bengio, Y., & Haffner, P. (1998). Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11), 2278–2324.
- LeCun, Y., Denker, J.S., & Solla, S.A. (1990). Optimal brain damage. In *Advances in neural information processing systems* (pp. 598–605).
- Lebedev, V., & Lempitsky, V. (2016). Fast convnets using group-wise brain damage. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 2554–2564).
- Lee, N., Ajanthan, T. & Torr, P.H. (2019). Snip: Single-shot network pruning based on connection sensitivity. In *International conference on learning representations (ICLR)*
- Li, H., Kadav, A., Durdanovic, I., Samet, H., & Graf, H.P. (2017). Pruning filters for efficient convnets. In *international conference of learning representation (ICLR)*.
- Li, T., Wu, B., Yang, Y., Fan, Y., Zhang, Y., & Liu, W. (2019). Compressing convolutional neural networks via factorized convolutional filters. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 3977–3986).
- Liao, N., Wang, S., Xiang, L., Ye, N., Shao, S., & Chu, P. (2022). Achieving adversarial robustness via sparsity. *Machine Learning*, 111(2), 685–711.

- Lin, M., Ji, R., Wang, Y., Zhang, Y., Zhang, B., Tian, Y., & Shao, L. (2020). Hrank: Filter pruning using high-rank feature map. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 1529–1538).
- Louizos, C., Ullrich, K., Welling, M. (2017). Bayesian compression for deep learning. In *Advances in neural information processing systems* (pp. 3288–3298).
- Luo, J.-H., Wu, J., & Lin, W. (2017). Thinet: A filter level pruning method for deep neural network compression. In *Proceedings of the IEEE international conference on computer vision* (pp. 5058–5066)
- Mallya, A., & Lazebnik, S. (2018). Packnet: Adding multiple tasks to a single network by iterative pruning. In: *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 7765–7773).
- Mallya, A., Davis, D., & Lazebnik, S. (2018) Piggyback: Adapting a single network to multiple tasks by learning to mask weights. In *Proceedings of the European conference on computer vision (ECCV)* (pp. 67–82).
- Mehta, D., Kim, K.I., & Theobalt, C. (2019). On implicit filter level sparsity in convolutional neural networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 520–528).
- Molchanov, P., Mallya, A., Tyree, S., Frosio, I., & Kautz, J. (2019). Importance estimation for neural network pruning. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 11264–11272).
- Molchanov, P., Tyree, S., Karras, T., Aila, T., & Kautz, J. (2016). Pruning convolutional neural networks for resource efficient inference. arXiv preprint [arXiv:1611.06440](https://arxiv.org/abs/1611.06440)
- Molchanov, D., Ashukha, A., & Vetrov, D. (2017). Variational dropout sparsifies deep neural networks. In *International conference on machine learning* (pp. 2498–2507). PMLR.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Köpf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J. & Chintala, S. (2019). Pytorch: An imperative style, high-performance deep learning library. arXiv preprint [arXiv:1912.01703](https://arxiv.org/abs/1912.01703)
- Renda, A., Frankle, J., & Carbin, M. (2020). Comparing rewinding and fine-tuning in neural network pruning. arXiv preprint [arXiv:2003.02389](https://arxiv.org/abs/2003.02389)
- Simonyan, K., & Zisserman, A. (2015). Very deep convolutional networks for large-scale image recognition. In *International conference on learning representations (ICLR)*.
- Sui, Y., Yin, M., Xie, Y., Phan, H., Aliari Zonouz, S., & Yuan, B. (2021). Chip: Channel independence-based pruning for compact neural networks. *Advances in Neural Information Processing Systems*, 34, 24604–24616.
- Ullrich, K., Meeds, E., & Welling, M. (2017). Soft weight-sharing for neural network compression. arXiv preprint [arXiv:1702.04008](https://arxiv.org/abs/1702.04008)
- Ye, S., Xu, K., Liu, S., Cheng, H., Lambrechts, J.-H., Zhang, H., Zhou, A., Ma, K., Wang, Y., & Lin, X. (2019). Adversarial robustness vs. model compression, or both?. In *Proceedings of the IEEE/CVF international conference on computer vision* (pp. 111–120).
- Ye, J., Lu, X., Lin, Z., & Wang, J.Z. (2018). Rethinking the smaller-norm-less-informative assumption in channel pruning of convolution layers. arXiv preprint [arXiv:1802.00124](https://arxiv.org/abs/1802.00124)
- Yu, R., Li, A., Chen, C.-F., Lai, J.-H., Morariu, V.I., Han, X., Gao, M., Lin, C.-Y., & Davis, L.S. (2018). Nisp: Pruning networks using neuron importance score propagation. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 9194–9203).
- Zagoruyko, S. (2015). 92.45 on cifar-10 in torch. URL: <http://torch.ch/blog/2015/07/30/cifar.html>.
- Zhao, C., Ni, B., Zhang, J., Zhao, Q., Zhang, W., & Tian, Q. (2019). Variational convolutional neural network pruning. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 2780–2789).

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.