



HAL
open science

An experimental study of the impact of pre-training on the pruning of a convolutional neural network

Nathan Hubens, Matei Mancas, Bernard Gosselin, Marius Preda, Titus Zaharia

► **To cite this version:**

Nathan Hubens, Matei Mancas, Bernard Gosselin, Marius Preda, Titus Zaharia. An experimental study of the impact of pre-training on the pruning of a convolutional neural network. 3rd International Conference on Applications of Intelligent Systems (APPIS), Jan 2020, Las Palmas de Gran Canaria Spain, France. pp.1-6, 10.1145/3378184.3378224 . hal-04389265

HAL Id: hal-04389265

<https://hal.science/hal-04389265v1>

Submitted on 27 Jan 2025

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.

AN EXPERIMENTAL STUDY OF THE IMPACT OF PRE-TRAINING ON THE PRUNING OF A CONVOLUTIONAL NEURAL NETWORK

Nathan Hubens ^{*†} Matei Mancas ^{*} Bernard Gosselin ^{*} Marius Preda [†] Titus Zaharia [†]

^{*} ISIA Lab, University of Mons

[†] Artemis, IP Paris

ABSTRACT

In recent years, deep neural networks have known a wide success in various application domains. However, they require important computational and memory resources, which severely hinders their deployment, notably on mobile devices or for real-time applications. Neural networks usually involve a large number of parameters, which correspond to the weights of the network. Such parameters, obtained with the help of a training process, are determinant for the performance of the network. However, they are also highly redundant. The pruning methods notably attempt to reduce the size of the parameter set, by identifying and removing the irrelevant weights. In this paper, we examine the impact of the training strategy on the pruning efficiency. Two training modalities are considered and compared: (1) fine-tuned and (2) from scratch. The experimental results obtained on four datasets (CIFAR10, CIFAR100, SVHN and Caltech101) and for two different CNNs (VGG16 and MobileNet) demonstrate that a network that has been pre-trained on a large corpus (e.g. ImageNet) and then fine-tuned on a particular dataset can be pruned much more efficiently (up to 80% of parameter reduction) than the same network trained from scratch.

Index Terms—Neural Network Pruning, Fine-tuning, CNN compression

1. INTRODUCTION

Over the last years, Convolutional Neural Networks (CNNs) have exhibited state-of-the-art performance in various computer vision tasks, including image classification and object detection [1, 2]. While their performance continuously improved, such networks kept growing deeper in complexity and depth. This led to increasing needs of parameters storage and convolution operations, inhibiting the use of deep neural networks for applications with memory or processing time limitations.

As the majority of the parameters of CNNs are located in the fully-connected layers, recent works have attempted to

reduce the memory demands of such networks by removing the fully-connected part of the network, which is replaced by an average pooling layer [3, 4]. Such approaches make it possible to decrease significantly the number of parameters without affecting the accuracy. On the other hand, the convolution layers are those where most of the computations are realized. This computation cost can be reduced by down-sampling the feature maps earlier in the network [5] through various pooling mechanisms or by replacing the convolution operations by factorized convolutions [6, 7].

Nowadays, a widely used approach when training neural networks for specific applications concerns the fine-tuning [8]. It consists of using the weights of a model pre-trained on a generic database as initialization and then training the model on a particular dataset. Neural network fine-tuning offers the advantage of making possible to obtain results for particular applications, where the training datasets are not sufficiently large. The dataset used for the pre-training is often very large and contains various examples (e.g. ImageNet [9]), making the network able to extract a large variety of different features. Generally, the network that has been pre-trained on the large dataset also has a very high capacity and can be over-parameterized for the particular dataset. This can lead to redundant or irrelevant weights. Removing such weights will then decrease the number of parameters of the network without significantly degrading its accuracy.

In this work, we propose an experimental analysis of the sensitivity of the pruning methods with respect to the training strategy. We consider for evaluation: (1) a network that has been pre-trained on a generic dataset (i.e., ImageNet), and then fine-tuned on a target dataset; and (2) a network that has been directly trained on the target dataset from randomly initialized weights. The experimental results obtained on various datasets (CIFAR10, CIFAR100, SVHN and Caltech101) and for two different networks (VGG16 and MobileNet) demonstrate the superiority of the first approach, which consists of using pre-training process.

The rest of the paper is organized as follows. Section 2 briefly presents an overview of the state-of-the-art prun-

¹This research has been conducted in the context of a joint-PhD between the two institutions.

ing methods. Section 3 introduces the retained evaluation methodology, with adopted pruning technique, datasets and network architectures. The experimental results obtained are presented and discussed in Section 4. Finally, Section 5 concludes the paper and opens perspectives of future work.

2. RELATED WORK

Recently, there has been a line of work concerning the lottery ticket hypothesis [10, 11, 12]. This hypothesis argues that, in a regular neural network architecture, there exists a sub-network that can be trained to the same level of performance as the original one, as long as it starts from the original initial conditions. This means that a neural network does not require all of its parameters to perform correctly and that having an over-parametrized neural network is only useful to find the "winning ticket". In practice, those sub-networks can be found by training the original network to convergence, removing the unnecessary weights, then resetting the value of the remaining weights to their original value.

Early work on pruning methods dates back to Optimal Brain Damage [13] and Optimal Brain Surgeon [14], where the weights are pruned based on the Hessian of the loss function. More recently, [15] proposes to prune the weights with small magnitude. This kind of pruning, performed on individual weights, is called unstructured pruning [16], as there is no intent to preserve any structure or geometry in the network. The unstructured pruning methods are the most efficient as they prune weights at the most fine-grained level. However, they lead to sparse weight matrices. Consequently, taking the advantage of the pruning results requires dedicated hardware or libraries able to efficiently deal with such sparsity.

To overcome this limitation, so-called structured pruning methods have been introduced. In this case, the pruning is operated at the level of the filters, kernels or vectors [17, 18]. One of the most popular structured pruning technique is the so-called filter pruning. Here, at each pruning stage, complete convolution filters are removed and the convolution structure remains unchanged. The filter selection methods can be based on the filter weight norm [19], average percentage of zeros in the output [20], or on the influence of each channel on the final loss [21]. Intrinsic correlation within each layer can also be exploited to remove redundant channels [22]. Pruning can also be combined with other compression techniques such as weight quantization [23], low-rank approximations of weights [24], knowledge distillation [25, 26] to further reduce the size of the network.

In our work, we have adopted a structured pruning approach, performed at the filter level. As in [19, 27], we use the l_1 -norm to determine the importance of filters and select those to prune. We decided to adopt an iterative pruning pro-

cess, where the pruning is performed one layer at a time, and the network is retrained after each pruning phase, to allow it to recover from the loss of some parameters.

3. METHODOLOGY

In this section, we describe in detail the methodology we followed for training and pruning a model.

3.1. Pruning method

The pruning method [19] used in our work prunes the filters of a trained network that have lowest sensitivity to pruning. The sensitivity to pruning of a layer is closely related to the l_1 -norm of its filters. Figure 1b shows the sensitivity to pruning of each layer of a VGG16 network fine-tuned on the MNIST dataset [28]. The layers that are the most sensitive, i.e. the layers where the accuracy drops the fastest when removing filters, are also the layers with the most high l_1 -norm filters, as illustrated in Figure 1a.

Based on this observation, we propose to remove filters with the lowest l_1 -norm as they produce feature maps with weaker activations compared to other filters in that layer.

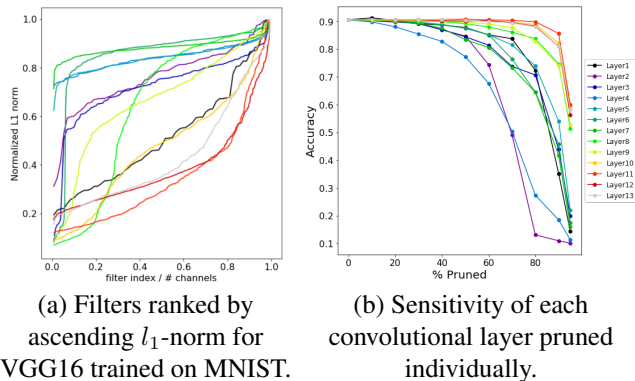


Fig. 1. Visualization of the sensitivity to pruning of a VGG-16 trained on MNIST.

More precisely, the procedure of pruning m filters from a layer i is the following:

1. Compute the l_1 -norm of each filter of layer i .
2. Sort the filters by their l_1 -norm.
3. Prune m filters with the smallest l_1 -norm and their corresponding feature maps in layer $i + 1$. The removed feature maps will not participate to the next computations, thus their corresponding kernels in the successive convolution filters are also removed (Figure 2).
4. Retrain the whole network until new convergence.

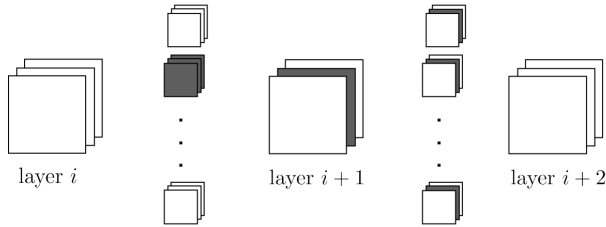


Fig. 2. Representation of filter pruning of a convolutional layer. The pruned filter, as well as its corresponding feature map and kernels in the following layer, are removed.

3.2. Datasets

We have carried out our experiments on four different datasets, with a good variability of number of classes, image resolution and content. As targeted datasets, we have CIFAR-10 and CIFAR-100 [29], consisting of RGB images of 32×32 pixels, labelled respectively over 10 and 100 classes. We also retained SVHN [30], a real-world image dataset for recognizing digits and numbers in natural scene images of size 32×32 . Finally, we have considered the Caltech101 corpus [31], consisting of pictures of objects belonging to 101 categories, of size of approximately 300×200 pixels.

3.3. Network architectures

In order to validate our results, we have adopted two well-known network architectures. The first one is VGG16 [1]. Here, we have replaced the original fully-connected layers by a Global Average Pooling layer and 2 narrow fully-connected layers. In this way, most parameters are contained in the convolutional layers. The network thus consists of 13 convolutional layers and 2 fully-connected layers.

The second network retained is MobileNetV1 [6], specifically designed to achieve efficiency both in parameter number and in computation complexity. MobileNet uses a factorized form of convolutions called Depthwise Separable Convolutions. The MobileNet architecture used in our experiments thus consists of one standard convolution layer acting on the input image, 13 depthwise separable convolutions, and finally a global average pooling and 2 fully connected layers.

Whatever the network retained, in experiments, Network-A will refer to the network pre-trained on ImageNet and fine-tuned on the target dataset, while Network-B refers to the same network trained on the target datasets from scratch.

4. EXPERIMENTAL RESULTS

In this section, we compare the pruning efficiency of both fine-tuned and trained from scratch networks. Our models are first trained until convergence, using the Adam optimizer, with an initial learning rate of 0.001 and a step decay scheduling, reducing the learning rate by a factor 10 every 40 epochs. They are then tested on the validation set to get the baseline accuracy.

After each pruning phase, a retraining is performed during 5 epochs, with the lowest learning rate reached during baseline training. We also monitor the accuracy on the validation set, and proceed to another pruning phase if it has not dropped by more than 1% from the baseline accuracy.

4.1. VGG16

To select the layer that will be pruned and decide of many filters to be removed, we conduct a sensitivity analysis on the network. As shown in Figure 3a, most of the low-norm filters are contained in the later layers. This suggests that those layers will be less sensitive to pruning than the others.

On the other hand, Figure 3b shows that Network-B has a more even repartition of filter norms, which suggests that it is more sensitive to pruning than Network-A, specifically in the later layers.

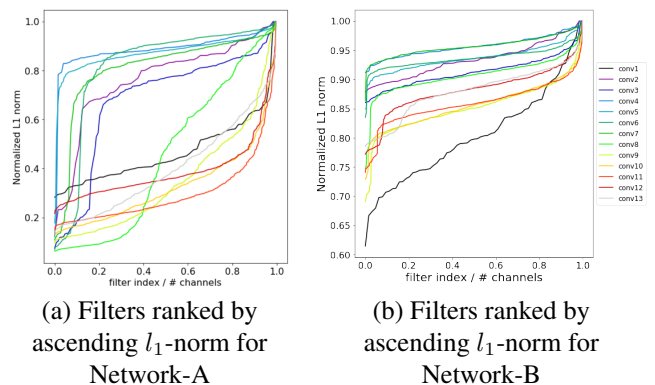


Fig. 3. Visualization of the importance of filters of VGG-16 trained on CIFAR-10. Filters are ranked by l_1 -norm.

This observation is confirmed by the results of pruning, summarized in Table 1, which shows that indeed, most of the pruning of Network-A can be performed in the later layers while pruning of Network-B is more evenly distributed.

Table 2 summarizes the results on all the tested datasets, and shows the resulting number of parameters, their corresponding storage size, as well as the number of floating point operation (FLOPs) needed for an input image to traverse the whole network at testing phase. We can observe that for all

Table 1. Parameters remaining for each layer after the pruning of VGG16, trained on CIFAR10.

Layer type	#Params	Network-A	Network-B
Conv 1	1792	1792	1792
Conv 2	36,928	36,928	36,928
Conv 3	73,856	73,856	73,856
Conv 4	147,584	147,584	129,136
Conv 5	295,168	295,168	161,440
Conv 6	590,080	590,080	230,560
Conv 7	590,080	590,080	230,560
Conv 8	1,180,160	442,560	553,344
Conv 9	2,359,808	331,968	1,327,488
Conv 10	2,359,808	331,968	1,327,488
Conv 11	2,359,808	221,312	1,327,488
Conv 12	2,359,808	147,584	1,327,488
Conv 13	2,359,808	147,584	1,327,488
Linear	262,656	66,048	197,120
Linear	5130	5130	5130
Total	14.98M	3.43M	8.26M

the tested datasets, pruning is more effective in terms of parameters removed for Network-A than for Network-B.

A second observation is that a lower number of parameters do not necessarily lead to a reduction in FLOPs. This phenomenon can be explained by the fact that, while most of the parameters are contained in the later layers, most of the operations are performed in the first ones, where the resolutions of the activation maps are higher. For Network-A, most of the pruning is performed in the later layers. In contrast, for Network-B, the pruning is more distributed throughout the network. Thus, Network-A has the fewest parameters but Network-B often has the fewest FLOPs.

Table 2. Results of the pruning on VGG16 for the four studied datasets.

Dataset	Network	Params (M)	FLOP (M)	Size (MB)
CIFAR10	Baseline	14.98	627.48	57.22
	A-pruned	3.43	421.20	13.15
	B-pruned	8.26	397.85	31.57
CIFAR100	Baseline	15.03	627.57	57.39
	A-pruned	8.26	503.47	31.57
	B-pruned	8.69	444.56	33.24
SVHN	Baseline	14.98	627.48	57.22
	A-pruned	3.15	414.12	12.10
	B-pruned	4.08	311.04	15.61
Caltech101	Baseline	15.03	30,720.99	57.40
	A-pruned	8.44	25,402.11	32.28
	B-pruned	13.59	30,171.17	51.93

A closer look at the filters of the first convolutional layer (Figure 4) exhibits the difference of learned filters. The reason of this difference is because the pre-training of Network-

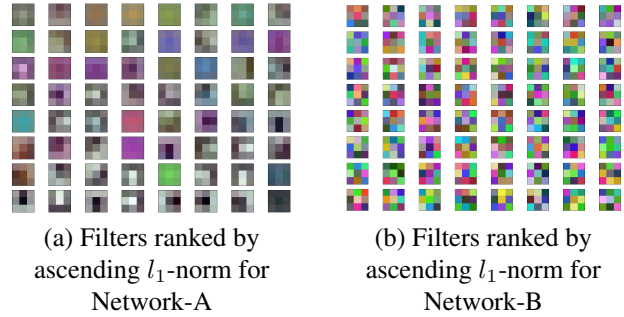


Fig. 4. Visualization of the 64 filters in the first convolutional layer of VGG-16 trained on CIFAR-10. Network-A filters have more structure than those of Network-B.

A helped to find useful filters on the ImageNet database. As Network-B only had access to fewer and smaller images, it could not learn the same kind of filters by itself and had to distribute the feature extraction across the network, reason why the later layers are more sensitive to pruning.

4.2. MobileNet

The Depthwise Separable Convolutions operations that are used in MobileNet are composed of two operations. The first is a Depthwise Convolution, that filters each input map independently. The second operation is called a Pointwise Convolution, that combines the results of the previous operation. The Pointwise Convolution is similar to a regular convolution but using filter dimensions of 1×1 . As most of the parameters are contained in the second operation, we decided to operate the pruning only on Pointwise Convolutions.

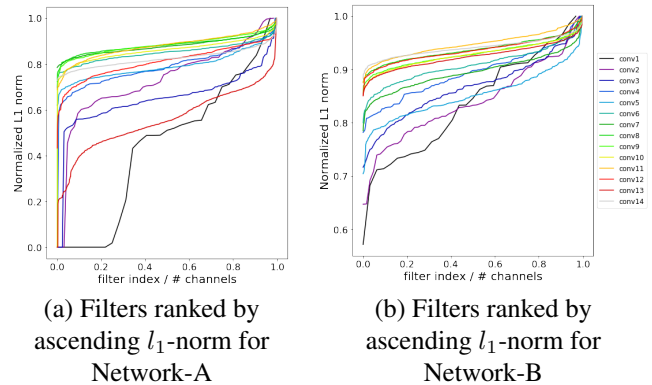


Fig. 5. Visualization of the importance of filters of MobileNet trained on CIFAR-10. Filters are ranked by l_1 -norm.

We follow the same process as for VGG16 and perform a sensitivity analysis on MobileNet before deciding which layer and how many filters to remove. The difference between the sensitivity of Network-A (Figure 5a) and Network-B (Figure 5b) is less clear than in the case of VGG16 but some useful

information can still be extracted. The highest norm filters of Network-B are still in the later layers, while it isn’t necessarily the case for Network-A. This again suggests that Network-A can be pruned further in the later layers, and confirmed in Table 3.

Table 3. Parameters remaining for each layer after the pruning of MobileNet, trained on CIFAR10. For clarity, only the first convolution and pointwise convolutions are represented.

Layer type	#Params	Network-A	Network-B
Conv 1	864	567	864
Conv 2	2048	1344	2048
Conv 3	8192	8192	8192
Conv 4	16,384	16,384	16,384
Conv 5	32,768	32,768	28,672
Conv 6	65,536	65,536	50,176
Conv 7	131,072	98,304	64,512
Conv 8	262,144	147,456	82,944
Conv 9	262,144	147,456	82,944
Conv 10	262,144	147,456	82,944
Conv 11	262,144	98,304	82,944
Conv 12	262,144	65,536	78,336
Conv 13	524,288	65,536	82,688
Conv 14	1,048,576	32,768	97,280
Linear	524,800	66,048	164,352
Linear	5130	5130	5130
Total	3.76M	1.05M	0.98M

The results of the pruning of MobileNet on the different datasets are summarized in Table 4. As it was the case for VGG16, Network-A can also be pruned further than Network-B, even if the difference is smaller in this case.

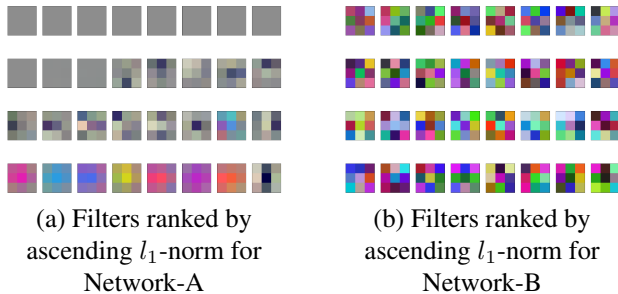


Fig. 6. Visualization of the 32 filters in the first convolutional layer of MobileNet trained on CIFAR-10. Filters of Network-A exhibit more structure than the filters of Network-B.

Again, looking at the filters of the first convolutional layer exhibits the difference of learned filters between Network-A (Figure 6a) and Network-B (Figure 6b). The filters of Network-A show some structure where those of Network-B don’t appear to. Moreover, as it was suggested in Figure 5a, Network-A has some really low-norm filters, that don’t seem to extract useful information, and can thus be removed.

Table 4. Results of the pruning on MobileNet for the four studied datasets.

Dataset	Network	Params (M)	FLOP (M)	Size (MB)
CIFAR10	Baseline	3.76	24.23	14.58
	A-pruned	1.05	13.83	4.26
	B-pruned	0.98	12.26	3.99
CIFAR100	Baseline	3.80	24.32	14.76
	A-pruned	2.54	21.42	9.94
	B-pruned	3.12	21.89	12.14
SVHN	Baseline	3.76	24.23	14.58
	A-pruned	1.50	17.66	5.99
	B-pruned	1.75	16.57	6.94
Caltech101	Baseline	3.81	1136.59	14.76
	A-pruned	3.08	1078.30	11.99
	B-pruned	3.43	1105.83	13.32

5. CONCLUSION AND PERSPECTIVES

In this paper, we have investigated the impact of the training process on the number of relevant filters involved in a CNN. In particular, we have compared the sensitivity of filter pruning of a network that has been fine-tuned to the sensitivity of a network trained from scratch.

Experiments have been conducted on four different datasets (CIFAR10, CIFAR100, SVHN, Caltech101) and for two different network architectures (VGG16, MobileNet). Results have shown that a CNN that has been pre-trained and then fine-tuned on a target dataset is less sensitive to pruning and thus, can be pruned further than the same network trained from-scratch on the target dataset.

This also supposes that the methodology of training has a strong impact on the part of the network where features are extracted. Pre-training helps the network to be discriminant in early layers and thus allows for more pruning in later layers, where most of the parameters are contained. On the other hand, training a network from randomly initialized weights makes its layers more evenly discriminant and thus the pruning has to be more distributed across the layers, resulting in fewer parameters removed but often a bigger reduction in computation operations.

Further research can be conducted on using this combination of sensitivity analysis and pruning to improve existing architectures of CNNs or find new ones that are more efficient. It can also lead to finding new training strategies either targeting a low number of parameters after pruning, as it is the case with fine-tuning, or targeting fewer computations, as it is the case for a network trained from-scratch.

6. REFERENCES

- [1] Karen Simonyan and Andrew Zisserman, “Very deep convolutional networks for large-scale image recognition,” in *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015.
- [2] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich, “Going deeper with convolutions,” in *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2015, pp. 1–9, ISSN: 1063-6919, 1063-6919.
- [3] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jonathon Shlens, and Zbigniew Wojna, “Rethinking the inception architecture for computer vision,” in *Conference on Computer Vision and Pattern Recognition, CVPR*, 2016.
- [4] Min Lin, Qiang Chen, and Shuicheng Yan, “Network in network,” in *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings*, 2014.
- [5] Kaiming He and Jian Sun, “Convolutional neural networks at constrained time cost,” *CoRR*, vol. abs/1412.1710, 2014.
- [6] Andrew G. Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam, “Mobilenets: Efficient convolutional neural networks for mobile vision applications,” *CoRR*, vol. abs/1704.04861, 2017.
- [7] François Chollet, “Xception: Deep learning with depthwise separable convolutions,” *CoRR*, vol. abs/1610.02357, 2016.
- [8] Jason Yosinski, Jeff Clune, Yoshua Bengio, and Hod Lipson, “How transferable are features in deep neural networks?,” *CoRR*, vol. abs/1411.1792, 2014.
- [9] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei, “ImageNet: A large-scale hierarchical image database,” in *2009 IEEE Conference on Computer Vision and Pattern Recognition*, 2009, pp. 248–255.
- [10] Jonathan Frankle and Michael Carbin, “The lottery ticket hypothesis: Training pruned neural networks,” *CoRR*, vol. abs/1803.03635, 2018.
- [11] Jonathan Frankle, Gintare Karolina Dziugaite, Daniel M. Roy, and Michael Carbin, “The lottery ticket hypothesis at scale,” *CoRR*, vol. abs/1903.01611, 2019.
- [12] Ari S. Morcos, Haonan Yu, Michela Paganini, and Yuandong Tian, “One ticket to win them all: generalizing lottery ticket initializations across datasets and optimizers,” *ArXiv*, vol. abs/1906.02773, 2019.
- [13] Yann LeCun, John S. Denker, and Sara A. Solla, “Optimal brain damage,” in *Advances in Neural Information Processing Systems 2*, D. S. Touretzky, Ed., pp. 598–605. Morgan-Kaufmann, 1990.
- [14] Babak Hassibi, David G. Stork, Gregory Wolff, and Takahiro Watanabe, “Optimal brain surgeon: Extensions and performance comparisons,” in *Advances in Neural Information Processing Systems 6*, pp. 263–270. Morgan-Kaufmann, 1994.
- [15] Song Han, Jeff Pool, John Tran, and William J. Dally, “Learning both weights and connections for efficient neural networks,” *CoRR*, vol. abs/1506.02626, 2015.
- [16] Sajid Anwar, Kyuyeon Hwang, and Wonyong Sung, “Structured pruning of deep convolutional neural networks,” *CoRR*, vol. abs/1512.08571, 2015.
- [17] Jian Cheng, Peisong Wang, Gang Li, Qinghao Hu, and Hanqing Lu, “Recent advances in efficient computation of deep convolutional neural networks,” *CoRR*, vol. abs/1802.00939, 2018.
- [18] Yihui He, Xiangyu Zhang, and Jian Sun, “Channel pruning for accelerating very deep neural networks,” *CoRR*, vol. abs/1707.06168, 2017.
- [19] Hao Li, Asim Kadav, Igor Durdanovic, Hanan Samet, and Hans Peter Graf, “Pruning filters for efficient convnets,” *CoRR*, vol. abs/1608.08710, 2016.
- [20] Hengyuan Hu, Rui Peng, Yu-Wing Tai, and Chi-Keung Tang, “Network trimming: A data-driven neuron pruning approach towards efficient deep architectures,” *CoRR*, vol. abs/1607.03250, 2016.
- [21] Pavlo Molchanov, Stephen Tyree, Tero Karras, Timo Aila, and Jan Kautz, “Pruning convolutional neural networks for resource efficient transfer learning,” *CoRR*, vol. abs/1611.06440, 2016.
- [22] Xavier Suau, Luca Zappella, Vinay Palakkode, and Nicholas Apostoloff, “Principal filter analysis for guided network compression,” *CoRR*, vol. abs/1807.10585, 2018.
- [23] Mohammad Rastegari, Vicente Ordonez, Joseph Redmon, and Ali Farhadi, “Xnor-net: Imagenet classification using binary convolutional neural networks,” *CoRR*, vol. abs/1603.05279, 2016.

- [24] Emily Denton, Wojciech Zaremba, Joan Bruna, Yann LeCun, and Rob Fergus, “Exploiting linear structure within convolutional networks for efficient evaluation,” *CoRR*, vol. abs/1404.0736, 2014.
- [25] Geoffrey Hinton, Oriol Vinyals, and Jeffrey Dean, “Distilling the knowledge in a neural network,” in *NIPS Deep Learning and Representation Learning Workshop*, 2015.
- [26] Jimmy Ba and Rich Caruana, “Do deep nets really need to be deep?,” in *Advances in Neural Information Processing Systems*, pp. 2654–2662. Curran Associates, Inc., 2014.
- [27] Zhuang Liu, Mingjie Sun, Tinghui Zhou, Gao Huang, and Trevor Darrell, “Rethinking the value of network pruning,” *CoRR*, vol. abs/1810.05270, 2018.
- [28] Yann LeCun and Corinna Cortes, “MNIST handwritten digit database,” 2010.
- [29] Alex Krizhevsky, “Learning multiple layers of features from tiny images,” 2009.
- [30] Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bis-sacco, Bo Wu, and Andrew Y. Ng, “Reading digits in natural images with unsupervised feature learning,” in *Workshop on Deep Learning and Unsupervised Feature Learning, NeurIPS*, 2011.
- [31] Fei-Fei Li, Rob Fergus, and Pietro Perona, “Learning generative visual models from few training examples: An incremental bayesian approach tested on 101 object categories,” in *2004 Conference on Computer Vision and Pattern Recognition Workshop*, June 2004, pp. 178–178.