

AperTO - Archivio Istituzionale Open Access dell'Università di Torino

Pruning Artificial Neural Networks: A Way to Find Well-Generalizing, High-Entropy Sharp Minima

This is the author's manuscript

Original Citation:

Availability:

This version is available <http://hdl.handle.net/2318/1765267> since 2020-12-28T11:14:23Z

Publisher:

Springer Science and Business Media Deutschland GmbH

Published version:

DOI:10.1007/978-3-030-61616-8_6

Terms of use:

Open Access

Anyone can freely access the full text of works made available as "Open Access". Works made available under a Creative Commons license can be used according to the terms and conditions of said license. Use of all other works requires consent of the right holder (author or publisher) if not exempted from copyright protection by the applicable law.

(Article begins on next page)

Pruning artificial neural networks: a way to find well-generalizing, high-entropy sharp minima

Enzo Tartaglione^[0000-0003-4274-8298], Andrea Bragagnolo, and
Marco Grangetto^[0000-0002-2709-7864]

University of Torino, Computer Science dept., Torino TO 10149, Italy
{enzo.tartaglione, andrea.bragagnolo}@unito.it

Abstract. Recently, a race towards the simplification of deep networks has begun, showing that it is effectively possible to reduce the size of these models with minimal or no performance loss. However, there is a general lack in understanding why these pruning strategies are effective. In this work, we are going to compare and analyze pruned solutions with two different pruning approaches, one-shot and gradual, showing the higher effectiveness of the latter. In particular, we find that gradual pruning allows access to narrow, well-generalizing minima, which are typically ignored when using one-shot approaches. In this work we also propose PSP-entropy, a measure to understand how a given neuron correlates to some specific learned classes. Interestingly, we observe that the features extracted by iteratively-pruned models are less correlated to specific classes, potentially making these models a better fit in transfer learning approaches.

Keywords: Pruning · Sharp minima · Entropy · Post synaptic potential · Deep learning.

1 Introduction

Artificial neural networks (ANNs) are nowadays one of the most studied algorithm used to solve a huge variety of tasks. Their success comes from their ability to learn from examples, not requiring any specific expertise and using very general learning strategies. The use of GPUs (and, recently, TPUs) for training ANNs gave a decisive kick to their large-scale deploy.

However, many deep models share a common obstacle: the large number of parameters, which allows their successful training [1,4], determines in turn a large number of operations at inference time, preventing efficient deployment to mobile and cheap embedded devices.

In order to address this problem, a number of approaches have been proposed, like defining new, more efficient models [10]. Recently, a race to shrink the size of these ANN models has begun: the so-called *pruning strategies* are indeed able to remove (or *prune*) non-relevant parameters from pre-trained models, reducing

the size of the ANN model, yet keeping a high generalization capability. On this topic, a very large amount of strategies have been proposed [6,14,19,21] from which we can identify two main classes:

- one-shot strategies: parameters are pruned using very fast, greedy approaches;
- gradual strategies: slower, potentially they can achieve higher *compression rates* (or in other words, they promise to prune more parameters at the cost of higher computational complexity).

In such a rush, however, an effort into a deeper understanding on potential properties of such sparse architectures has been mostly set aside: is there a specific reason for which we are able to prune many parameters with minimal or no generalization loss? Are one-shot strategies enough to match gradual pruning approaches? Is there any hidden property behind these sparse architectures?

In this work, we first compare one-shot pruning strategies to their gradual counterparts, investigating the benefits of having a much more computationally-intensive sparsifying strategy. Then, we shine a light on some local properties of minima achieved using the two different pruning strategies and finally, we propose *PSP-entropy*, a measure on the state of ReLU-activated neurons, to be used as an analysis tool to get a better understanding for the obtained sparse ANN models.

The rest of this paper is organized as follows. Sec. 2 reviews the importance of network pruning and the most relevant literature. Next, in Sec. 3 we discuss the relevant literature around local properties of minima for ANN models. Then, in Sec. 4 we propose PSP-entropy, a metric to measure how much a neuron specializes in identifying features belonging to a sub-set of classes learned at training time. Sec. 5 provides our findings on the properties for sparse architectures and finally, in Sec. 6, we draw the conclusions and identify further directions for future research.

2 State of the art pruning techniques

In the literature it is possible to find a large number of pruning approaches, some old-fashioned [12] and others more recent [9,13,17]. Among the latter, many sub-categories can be identified. Ullrich *et al.* introduce what they call *soft weight sharing*, through which is possible to introduce redundancy in the network and reduce the amount of stored parameters [23]. Other approaches are based on parameters regularization and pruning: for example, Louizos *et al.* use an L_0 proxy regularization; Tartaglione *et al.*, instead, define the importance of a parameter via a sensitivity measure used as regularization [21]. Other approaches are dropout-based, like *sparse variational dropout*, proposed by Molchanov *et al.*, leveraging on a Bayesian interpretation of Gaussian dropout and promoting sparsity in the ANN model [17].

Overall, as stated in Sec. 1, most of the proposed pruning techniques can be divided in two macro classes. The first is defined by approaches based on *gradual pruning* [15,19,25], in which the network is, at the same time, trained and

pruned following some heuristic approach, spanning a large number of pruning iterations. One among these, showing the best performances, is LOBSTER, where parameters are gradually pruned according to their local contribution to the loss [19]. The second class, instead, includes all the techniques based on *one-shot pruning* [6,9,16]: here the pruning procedure consists of three stages:

1. a large, over-parametrized network is normally trained to completion;
2. the network is then pruned using some kind of heuristic (e.g. magnitude thresholding) to satisfaction (the percentage of remaining parameters is typically an hyper-parameter);
3. the pruned model is further fine-tuned to recover the accuracy lost due to the pruning stage.

A recent work by Frankle and Carbin [6] proposed the *lottery ticket hypothesis*, which is having a large impact on the research community. They claim that from an ANN, early in the training, it is possible to extract a sparse sub-network on a one-shot fashion: such sparse network, when trained, can match the accuracy of the original model. In a follow-up, Renda *et al.* propose a retraining approach that replaces the fine-tuning step with *weight rewinding*: after pruning, the remaining parameters are reset to their initial values and the pruned network is trained again. They also argue that using the initial weights values is fundamental to achieve competitive performance, which is degraded when starting from a random initialization [18].

On the other hand, Liu *et al.* show that, even when retraining a pruned sub-network using a new random initialization, they are able to reach an accuracy level comparable to its dense counterpart; challenging one of the conjectures proposed alongside the lottery ticket hypothesis [14].

3 Local properties of minima

In the previous section we have explored some of the most relevant pruning strategies. All of them rely on state-of-the-art optimization strategies: applying very simple optimizing heuristics to minimize the loss function, like for example SGD [2,26], it is nowadays possible to succeed in training ANNs on huge datasets. These problems are typically over-parametrized, and the dimensionality of the deep model trained can be efficiently reduced with almost no performance loss [21]. Furthermore, minimizing non-convex objective functions is typically supposed to make the trained architecture stuck into local minima. However, the empirical evidence shows that something else is happening under the hood: understanding it is in general of interest in order to improve the learning strategies.

Goodfellow *et al.* observed there is essentially no loss barrier between a generic random initialization for the ANN model and the final configuration [8]. Such a phenomena has also been observed on larger architectures by Draxler *et al.* [5]. These works lay as basis for the “lottery ticket hypothesis” papers. However, a

secondary yet relevant observation in [8] stated that there is a loss barrier between different ANN configurations showing similar generalization capabilities. Later, it was shown that typically a low loss path between well-generalizing solutions to the same learning problem can be found [20]. From this brief discussion it is evident that a general approach on how to better characterize such minima has yet to be found.

Keskar *et al.* showed why we should prefer small batch methods to large batch ones: they correlate the stochasticity introduced by small-batch methods to the sharpness of the reached minimum [11]. In general, they observe that the larger the batch, the sharper the reached minimum. Even more interestingly, they observe that the sharper the minimum, the worse the generalization of the ANN model. In general, there are many works supporting the hypothesis that flat minima generalize well, and this has been also the strength for a significant part of the current research [3,11]. However, in general this does not necessarily mean that no sharp minimum generalizes well, as we will see in Sec. 5.2.

4 Towards a deeper understanding: an entropy-based approach

In this section we propose PSP-entropy, a metric to evaluate the dependence of the output for a given neuron in the ANN model to the target classification task. The proposed measure will allow us to better understand the effect of pruning.

4.1 Post-synaptic potential

Let us define the output of the given i -th neuron at the l -th layer as

$$y_{l,i} = \varphi [f(\mathbf{y}_{l-1}, \theta_{l,i})] \quad (1)$$

where \mathbf{y}_{l-1} is the input of such neuron, $\theta_{l,i}$ are the parameters associated to it, $f(\cdot)$ is some affine function and $\varphi(\cdot)$ is the activation function, we can define its *post-synaptic potential* (PSP) [22] as

$$z_{l,i} = f(\mathbf{y}_{l-1}, \theta_{l,i}) \quad (2)$$

Typically, deep models are ReLU-activated: here on, let us consider the activation function for all the neurons in hidden layers as $\varphi(\cdot) = \text{ReLU}(\cdot)$. Under such assumption it is straightforward to identify two distinct regions for the neuron activation:

- $z_{l,i} \leq 0$: the output of the neuron will be exactly zero $\forall z_{l,i} \leq 0$;
- $z_{l,i} > 0$: there is a linear dependence of the output to $z_{l,i}$.

Hence, let us define

$$\varphi'(z) = \begin{cases} 0 & z \leq 0 \\ 1 & z > 0 \end{cases} \quad (3)$$

Intuitively, we understand that if two neurons belonging to the same layer, for the same input, share the same $\varphi'(z)$, then they are linearly-mappable to one equivalent neuron:

- $z_{l,i} \leq 0, z_{l,j} \leq 0$: one of them can be simply removed;
- $z_{l,i} > 0, z_{l,j} > 0$: they are equivalent to a linear combination of them.

In this work we are not interested in using this approach towards structured pruning: in the next section we are going to formulate a metric to evaluate the degree of disorder in the post synaptic potentials. The aim of such measure will be to have an analytical tool to give us a broader understanding on the behavior of the neurons in sparse architectures.

4.2 PSP-entropy for ReLU-activated neurons

In the previous section we have recalled the concept of post-synaptic potential. Some interesting concepts have been also introduced for ReLU-activated networks: we can use its value to approach the problem of *binning* the state of a neuron, according to $\varphi'(z_{l,i})$. Hence, we can construct a binary random process that we can rank according to its entropy. To this end, let us assume we set as input of our ANN model two different patterns, $\mu_{c,1}$ and $\mu_{c,2}$, belonging to the same class c (for those inputs, we aim at having the same target at the output of the ANN model). Let us consider the PSP $z_{l,i}$ (where l is a hidden layer):

- if $\varphi'(z_{l,i}|\mu_{c,1}) = \varphi'(z_{l,i}|\mu_{c,2})$ we can say there is *low PSP entropy*;
- if $\varphi'(z_{l,i}|\mu_{c,1}) \neq \varphi'(z_{l,i}|\mu_{c,2})$ we can say there is *high PSP entropy*.

We can model an entropy measure for PSP:

$$H(z_{l,i}|c) = - \sum_{t=\{0,1\}} p[\varphi'(z_{l,i}) = t|c] \cdot \log_2 \{p[\varphi'(z_{l,i}) = t|c]\} \quad (4)$$

where $p[\varphi'(z_{l,i}) = t|c]$ is the probability $\varphi'(z_{l,i}) = t$ when presented an input belonging to the c -th class. Since we typically aim at solving a multi-class problem, we can model an overall entropy for the neuron as

$$H(z_{l,i}) = \sum_c H(z_{l,i}|c) \quad (5)$$

It is very important to separate the contributions of the entropy according to the c -th target class since we expect the neurons to catch relevant features being highly-correlated to the target classes. Eq. (5) provides us very important information towards this end: the lower its value the more it specializes for some specific classes.

The formulation in (5) is very general and it can be easily extended to higher-order entropy, i.e. entropy of sets of neurons whose state correlates for the same classes. Now we are ready to use this metrics to shed further light to the findings in Sec. 5.

5 Experiments

For our test, we have decided to compare the state-of-the-art one-shot pruning proposed by Frankle and Carbin [6] to one of the top-performing gradual pruning strategies, LOBSTER [19]. Towards this end, we first obtain a sparse network model using LOBSTER; the non-pruned parameters are then re-initialized to their original values, according to the lottery ticket hypothesis [6]. Our purpose here is to determine whether the lottery ticket hypothesis applies also to the sparse models obtained using high-performing gradual pruning strategies.

As a second experiment, we want to test the effects of different, random initialization while keeping the achieved sparse architecture. According to Liu *et al.*, this should lead to similar results to those obtained with the original initialization [14]. Towards this end, we tried 10 different new starting configurations. As a last experiment, we want to assess how important is the structure originating from the pruning algorithm in reaching competitive performances after re-initialization: for this purpose, we randomly define a new pruned architecture with the same number of pruned parameters as those found via LOBSTER. Also in this case, 10 different structures have been tested.

We decided to experiment with different architectures and datasets commonly employed in the relevant literature: LeNet-300 and LeNet-5-caffe trained on MNIST, LeNet-5-caffe trained on Fashion-MNIST [24] and ResNet-32 trained on CIFAR-10.¹ For all our trainings we used the SGD optimization method with standard hyper-parameters and data augmentation, as defined in the papers of the different compared techniques [6,14,19].


5.1 One-shot vs gradual pruning

In Fig. 1 we show, for different percentages of pruned parameters, a comparison between the test accuracy of models pruned using the LOBSTER technique and the models retrained following the approaches we previously defined.


We can clearly identify a low compression rate regime in which the re-initialized model is able to recover the original accuracy, validating the lottery ticket hypothesis. On the other hand, when the compression rate rises (for example when we remove more than 95% of the LeNet-300 model’s parameters, as observed in Fig. 1a), the re-training approach strives in achieving low classification errors.

As one might expect, other combinations of dataset and models might react differently. For example, LeNet-300 is no longer able to reproduce the original performance when composed of less than 5% of the original parameters. On the other hand, LeNet-5, when applied on MNIST, is able to achieve an accuracy of around 99.20% even when 98% of its parameters are pruned away (Fig. 1b). This does not happen when applied on a more complex dataset like Fashion-MNIST, where removing 80% of the parameters already leads to performance degradation (Fig. 1c). Such a gap becomes extremely evident when we re-initialize an even more complex architecture like ResNet-32 trained on CIFAR-10 (Fig. 1d).


¹ https://github.com/akamaster/pytorch_resnet_cifar10



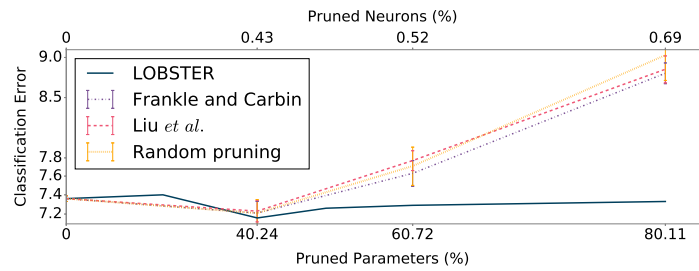
(a)



(b)



(c)



(d)

Fig.1: Test set error for different compression rates: LeNet-300 trained on MNIST (a), LeNet-5 trained on MNIST (b), LeNet-5 trained on Fashion-MNIST (c) and ResNet-32 trained on CIFAR-10 (d).

From the reported results, we observe that the original initialization is not always important: the error gap between a randomly initialized model and a model using the original weights’ values is minor, with the latter being slightly better. Furthermore, they both fail in recovering the performance for high compression rates.

5.2 Sharp minima can also generalize well




Fig. 2: Results of LeNet-5 trained on MNIST with the highest compression (99.57%): (a) plots the loss in the training set and (b) plots the magnitude of the top-5 largest hessian eigenvalues. G is the solution found with gradual learning while 1-S is the best one-shot solution (Frankle and Carbin).

In order to study the sharpness of local minima, let us focus, for example, on the results obtained on LeNet-5 trained on MNIST. We choose to focus our attention on this particular ANN model since, according to the state-of-the-art and coherently to our findings, we observe the lowest performance gap between gradual and one-shot pruning (as depicted in Fig. 1b); hence, it is a more challenging scenario to observe qualitative differences between the two approaches. However, we remark that all the observations for such a case apply also to the other architectures/datasets explored in Sec. 5.1.

In order to obtain the maps in Fig. 2, we follow the approach proposed by [8] and we plot the loss for the ANN configurations between two reference ones: in our case, we compare a solution found with gradual pruning (G) and one-shot (1-S). Then, we take a random orthogonal direction to generate a 2D map. Fig. 2a shows the loss on the training set between iterative and one-shot pruning for the highest compression rate (99.57% of pruned parameters as shown in Fig. 1b). According to our previous findings, we see that iterative pruning

lies in a lower loss region. Here, we show also the plot of the top-5 Hessian eigenvalues (all positive), in Fig. 2b, using the efficient approach proposed in [7]. Very interestingly, we observe that the solution proposed by iterative pruning lies in a narrower minimum than the one found using the one-shot strategy, despite generalizing slightly better. With this, we do not claim that narrower minima generalize well: gradual pruning strategies enable access to a *subset of well-generalizing narrow minima*, showing that not all the narrow minima generalize worse than the wide ones. This finding raises warnings against second order optimization, which might favor the research of flatter, wider minima, ignoring well-generalizing narrow minima. These non-trivial solutions are naturally found using gradual pruning which cannot be found using one-shot approaches, which on the contrary focus their effort on larger minima. In general, the sharpness of these minima explains why, for high compression rates, re-training strategies fail in recovering the performance, considering that it is in general harder to access this class of minima.

5.3 Study on the post synaptic potential




Fig. 3: L2 norm of PSP values for LeNet-5 trained on MNIST with 99.57% of pruned parameters.

In Sec. 5.2 we have observed that, as a result, iterative strategies focus on well-generalizing sharp minima. Is there something else yet to say about those? Let us inspect the average magnitude values of the PSPs for the different found solutions: towards this end, we could plot the average of their L2 norm values (z^2). As a first finding, gradually-pruned architectures naturally have lower PSP L2-norm values, as we observe in Fig. 3. None of the used pruning strategies explicitly minimize the term in z^2 : they naturally drive the learning towards such regions. However, the solution showing better generalization capabilities shows lower z^2 values. Of course, there are regions with even lower z^2 values; however, according to Fig. 2a, they should be excluded since they correspond to high-loss values (not all the low z^2 regions are low-loss). If we look at the PSP-entropy formulated in (5), we observe something interesting: gradual and




Fig. 4: Results on LeNet-5 trained on MNIST with 99.57% of pruned parameters. (a) plots the first order PSP-entropy, while (b) shows the second-order PSP entropy.

one-shot pruning show comparable first-order entropies, as shown in Fig. 4a.² It is interesting to see that there are also lower entropy regions which however correspond to higher loss values, according to Fig. 2a. When we move to higher-order entropies, something even more interesting happens: gradual pruning shows higher entropy than one-shot, as depicted in Fig. 4b (displaying the second order entropy). In such a case, having a lower entropy means having more groups of neurons specializing to specific patterns which correlate to the target class; on the contrary, having higher entropy yet showing better generalization performance results in having more general features, more agnostic towards a specific class, which still allow a correct classification performed by the output layer. This counter-intuitive finding has potentially-huge applications in transfer learning and domain adaptation, where it is critical to extract more general features, not very specific to the originally-trained problem.

6 Conclusion

In this work we have compared one-shot and gradual pruning on different state-of-the-art architectures and datasets. In particular, we have focused our attention in understanding potential differences and limits of both approaches towards achieving sparsity in ANN models.

We have observed that one-shot strategies are very efficient to achieve moderate sparsity at a lower computational cost. However, there is a limit to the maximum achievable sparsity, which can be overcome using gradual pruning.

² the source code for PSP-entropy is available at <https://github.com/EIDOSlab/PSP-entropy.git>

The highly-sparse architectures, interestingly, focus on a subset of sharp minima which are able to generalize well, which pose some questions to the potential sub-optimality of second-order optimization in such scenarios. This explains why we observe that one-shot strategies fail in recovering the performance for high compression rates. More importantly, we have observed, contrarily to what it could be expected, that highly-sparse gradually-pruned architectures are able to extract general features non-strictly correlated to the trained classes, making them unexpectedly, potentially, a good match for transfer-learning scenarios. Future works include a quantitative study on transfer-learning for sparse architectures and PSP-entropy maximization-based learning.

References

1. Ba, J., Caruana, R.: Do deep nets really need to be deep? In: Advances in neural information processing systems. pp. 2654–2662 (2014)
2. Bottou, L.: Large-scale machine learning with stochastic gradient descent. In: Proceedings of COMPSTAT’2010, pp. 177–186. Springer (2010)
3. Chaudhari, P., Choromanska, A., Soatto, S., LeCun, Y., Baldassi, C., Borgs, C., Chayes, J., Sagun, L., Zecchina, R.: Entropy-sgd: Biasing gradient descent into wide valleys. International Conference on Learning Representations, ICLR 2017 (2017)
4. Denton, E.L., Zaremba, W., Bruna, J., LeCun, Y., Fergus, R.: Exploiting linear structure within convolutional networks for efficient evaluation. In: Advances in neural information processing systems. pp. 1269–1277 (2014)
5. Draxler, F., Veschgini, K., Salmhofer, M., Hamprecht, F.: Essentially no barriers in neural network energy landscape. In: Dy, J., Krause, A. (eds.) Proceedings of the 35th International Conference on Machine Learning. Proceedings of Machine Learning Research, vol. 80, pp. 1309–1318. PMLR, Stockholmsmässan, Stockholm Sweden (10–15 Jul 2018), <http://proceedings.mlr.press/v80/draxler18a.html>
6. Frankle, J., Carbin, M.: The lottery ticket hypothesis: Finding sparse, trainable neural networks (2019), <https://www.scopus.com/inward/record.uri?eid=2-s2.0-85069453436&partnerID=40&md5=fd1a2b2384d79f66a49cc838a76343d3>, cited By 8
7. Golmant, N., Yao, Z., Gholami, A., Mahoney, M., Gonzalez, J.: pytorch-hessian-eigentings: efficient pytorch hessian eigendecomposition (Oct 2018), <https://github.com/noahgolmant/pytorch-hessian-eigentings>
8. Goodfellow, I.J., Vinyals, O., Saxe, A.M.: Qualitatively characterizing neural network optimization problems. International Conference on Learning Representations, ICLR 2015 (2015)
9. Han, S., Pool, J., Tran, J., Dally, W.: Learning both weights and connections for efficient neural network. In: Advances in neural information processing systems. pp. 1135–1143 (2015)
10. Howard, A.G., Zhu, M., Chen, B., Kalenichenko, D., Wang, W., Weyand, T., Andreetto, M., Adam, H.: Mobilenets: Efficient convolutional neural networks for mobile vision applications. arXiv preprint arXiv:1704.04861 (2017)
11. Keskar, N.S., Mudigere, D., Nocedal, J., Smelyanskiy, M., Tang, P.T.P.: On large-batch training for deep learning: Generalization gap and sharp minima. arXiv preprint arXiv:1609.04836 (2016)

12. LeCun, Y., Denker, J.S., Solla, S.A.: Optimal brain damage. In: *Advances in neural information processing systems*. pp. 598–605 (1990)
13. Li, H., Kadav, A., Durdanovic, I., Samet, H., Graf, H.P.: Pruning filters for efficient convnets. *International Conference on Learning Representations, ICLR 2017* (2017)
14. Liu, Z., Sun, M., Zhou, T., Huang, G., Darrell, T.: Rethinking the value of network pruning. *International Conference on Learning Representations, ICLR 2019* (2019)
15. Louizos, C., Welling, M., Kingma, D.P.: Learning sparse neural networks through l_0 regularization. *International Conference on Learning Representations, ICLR 2018* (2018)
16. Luo, J.H., Wu, J., Lin, W.: Thinet: A filter level pruning method for deep neural network compression. In: *Proceedings of the IEEE international conference on computer vision*. pp. 5058–5066 (2017)
17. Molchanov, D., Ashukha, A., Vetrov, D.: Variational dropout sparsifies deep neural networks. In: *34th International Conference on Machine Learning, ICML 2017*. vol. 5, pp. 3854–3863 (2017), <https://www.scopus.com/inward/record.uri?eid=2-s2.0-85048506601&partnerID=40&md5=c352a4786ef977ccea7e397bd7469f14>, cited By 29
18. Renda, A., Frankle, J., Carbin, M.: Comparing rewinding and fine-tuning in neural network pruning. *arXiv preprint arXiv:2003.02389* (2020)
19. Tartaglione, E., Bragagnolo, A., Grangetto, M., Lepsø y, S.: Loss-based sensitivity regularization:towards deep sparse neural networks. <https://iris.unito.it/retrieve/handle/2318/1737767/608158/ICML20.pdf> (2020)
20. Tartaglione, E., Grangetto, M.: Take a ramble into solution spaces for classification problems in neural networks. In: *International Conference on Image Analysis and Processing*. pp. 345–355. Springer (2019)
21. Tartaglione, E., Lepsø y, S., Fiandrotti, A., Francini, G.: Learning sparse neural networks via sensitivity-driven regularization. In: *Advances in Neural Information Processing Systems*. pp. 3878–3888 (2018)
22. Tartaglione, E., Perlo, D., Grangetto, M.: Post-synaptic potential regularization has potential. In: *International Conference on Artificial Neural Networks*. pp. 187–200. Springer (2019)
23. Ullrich, K., Welling, M., Meeds, E.: Soft weight-sharing for neural network compression. In: *5th International Conference on Learning Representations, ICLR 2017 - Conference Track Proceedings* (2019), <https://www.scopus.com/inward/record.uri?eid=2-s2.0-85071003624&partnerID=40&md5=dc00c36113f775ff4a6978b86543814d>, cited By 2
24. Xiao, H., Rasul, K., Vollgraf, R.: Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *CoRR* **abs/1708.07747** (2017), <http://arxiv.org/abs/1708.07747>
25. Zhu, M., Gupta, S.: To prune, or not to prune: exploring the efficacy of pruning for model compression. *International Conference on Learning Representations, ICLR 2018* (2018)
26. Zinkevich, M., Weimer, M., Li, L., Smola, A.J.: Parallelized stochastic gradient descent. In: *Advances in neural information processing systems*. pp. 2595–2603 (2010)