

# ONE-CYCLE PRUNING: PRUNING CONVNETS WITH TIGHT TRAINING BUDGET

Nathan Hubens <sup>\*†</sup>

Matei Mancaş <sup>\*</sup>

Bernard Gosselin <sup>\*</sup>

Marius Preda <sup>†</sup>

Titus Zaharia <sup>†</sup>

<sup>\*</sup> ISIA Lab, University of Mons, Belgium

<sup>†</sup> Artemis, IP Paris, France

## ABSTRACT

Introducing sparsity in a convnet has been an efficient way to reduce its complexity while keeping its performance almost intact. Most of the time, sparsity is introduced using a three-stage pipeline: 1) training the model to convergence, 2) pruning the model, 3) fine-tuning the pruned model to recover performance. The last two steps are often performed iteratively, leading to reasonable results but also to a time-consuming process. In our work, we propose to remove the first step of the pipeline and to combine the two others in a single training-pruning cycle, allowing the model to jointly learn the optimal weights while being pruned. We do this by introducing a novel pruning schedule, named One-Cycle Pruning (OCP), which starts pruning from the beginning of the training, and until its very end. Experiments conducted on a variety of combinations between architectures (VGG-16, ResNet-18), datasets (CIFAR-10, CIFAR-100, Caltech-101), and sparsity values (80%, 90%, 95%) show that not only OCP consistently outperforms common pruning schedules such as One-Shot, Iterative and Automated Gradual Pruning, but also that it drastically reduces the required training budget. Moreover, experiments following the Lottery Ticket Hypothesis show that OCP allows to find higher quality and more stable pruned networks.

**Index Terms**— Neural Network Pruning, Neural Network Compression, Pruning Schedule

## 1. INTRODUCTION

Deep neural networks are able to achieve state-of-the-art results in a wide variety of domains, including computer vision, natural language processing and speech recognition. But such achievements imply important increase in budget required both at training and inference time. More specifically, model size, run time memory and the number of computing operations are all constraints that make modern neural networks challenging to deploy on resource-constrained environments such as mobile phones or embedded devices.

For those reasons, neural network compression and acceleration have become active fields of research. Several re-

searchers have been interested in creating parameter-efficient architectures, by using low-rank approximations, parameters quantization, and neural network pruning.

Recent studies have also exhibited a particular characteristic of neural networks, called the *Lottery Ticket Hypothesis* [1], which suggests that, in regular architectures, there exists a sub-network that can be trained to at least the same level of performance as the original one, in a comparable training budget, as long as it starts from the same original conditions. Such a sub-network is said to have “won” at the initialization lottery and can be preserved, while other parameters of the original network can be removed using pruning methods.

To prune a neural network, the most commonly used technique is the so-called Iterative Pruning [2], requiring several cycles of pruning and fine-tuning which leads to a lengthy process. In this work, we propose to adopt a novel pruning schedule applied directly from the start of the initial training phase. This scheduling function gradually prunes the network during the training phase, thus making training and pruning of a neural network a joint process.

Contributions of our work are summarized as following:

- We propose a novel pruning schedule with stable, thus generic parameters.
- We show that the proposed pruning schedule, performed during the initial training of a neural network, drastically reduces the required training budget.
- We empirically show that our pruning schedule is able to find more stable lottery tickets, both for low and extreme sparsity levels, than other commonly used pruning schedules.

## 2. RELATED WORK

Pruning techniques can differ in many aspects. The main points of differentiation are presented in this section.

**Granularity.** The granularity used for the pruning is often categorized into two groups: *unstructured*, or when the pruning focuses on removing individual weights in the network and that there is no intent to keep any structure in the filters [3, 4]. Such a pruning method leads to sparse weight

---

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

matrices, requiring dedicated hardware or software to take advantage of the speed and computation gains. To overcome this limitation, *structured* pruning was introduced [5, 6, 7], which takes care of removing complete blocks of weights such as vectors, kernels, or even convolution filters.

**Criteria.** Early works on pruning criteria make use of second-order approximation of the loss surface to remove parameters [3, 4]. Other works have also explored the use of  $l_0$  regularization [8] during the training or even the use of variational dropout [9]. In addition to being more complex, those criteria often show to be less consistent across datasets and to lead to comparable or worse results than simple magnitude pruning, based on  $l_1$ -norm [10]. Moreover, the criteria can be used to compare weights belonging to the same layer, *i.e.* *local* pruning, leading to layers of equivalent sparsities. The criteria can also be applied to the whole network, comparing weights from all layers, *i.e.* *global* pruning, and leading to a network with layers of different sparsity levels.

**Scheduling.** While early pruning methods cared about removing redundant weights in a single step, called one-shot pruning [5], the most adopted technique nowadays is to perform pruning iteratively [11], starting from a pretrained network and performing several steps of pruning followed by fine-tuning. While having undeniably shown good results, such a pruning schedule is usually time consuming to obtain the final, pruned network [5]. For this reason, researchers have proposed alternatives to perform pruning during the training or even before it even started [12, 13]. Recently, some work has shown that the most critical phase of the training of a neural network happens during the very first iterations [14, 15] and that applying regularization after that initial phase has little effect on the final performance of the network [16]. One thus should apply pruning early in training to take advantage of its regularization effects but must do so very carefully to not irremediably damage the network during this brittle period [17, 18].

For those reasons, we propose a pruning schedule that is applied gradually during the training, and designed to be gentle during the very first iterations. This pruning schedule can be used for any granularity and any selection criteria but in this work, we are focusing on unstructured pruning, *i.e.* removing individual weights, and the criteria used for weights selection is their  $l_1$ -norms, compared locally in each layer.


### 3. THE PROPOSED SCHEDULE

The proposed method, One-Cycle Pruning (OCP), consists in starting from a dense network and inducing sparsity during the whole training phase. More precisely, sparsity is induced in the network according to the following schedule:

$$s_t = s_i + (s_f - s_i) \cdot \frac{1 + e^{-\alpha t + \beta}}{1 + e^{-\alpha + \beta}} \quad (1)$$

with  $s_t$ , the level of sparsity at training step  $t$ ,  $s_i$  and  $s_f$

respectively the initial and final level of sparsity, and  $\alpha$ ,  $\beta$  being two tuning parameters, modifying either the steepness of the scheduling (Figure 1a), or its horizontal offset (Figure 1b), to better suit the problem or architecture that is used.



**Fig. 1.** Visualization of the variation of the One-Cycle Pruning schedule for different  $\alpha$  and  $\beta$  values.

We performed a grid-search to find the best pair of  $\alpha$  and  $\beta$  values and found it to be  $\alpha = 14$  and  $\beta = 5$ . As reported in Table 1, OCP is stable around its optimal parameter values. They thus can be used as default, and being adapted if needed.

		$\beta$				
		3	4	5	6	7
$\alpha$	13	93.10 $\pm$ 0.18	93.23 $\pm$ 0.12	93.30 $\pm$ 0.07	93.31 $\pm$ 0.05	92.97 $\pm$ 0.07
	14	93.13 $\pm$ 0.06	93.16 $\pm$ 0.07	<b>93.46 <math>\pm</math> 0.13</b>	93.09 $\pm$ 0.18	93.25 $\pm$ 0.12
	15	93.10 $\pm$ 0.03	93.21 $\pm$ 0.14	93.19 $\pm$ 0.19	93.17 $\pm$ 0.08	93.28 $\pm$ 0.03

**Table 1.** Grid search of  $\alpha$  and  $\beta$  for Resnet-18 trained on CIFAR-10 for 90% sparsity. Mean and standard deviation of validation accuracy over 3 rounds are reported.

The benefits of such a schedule are three-fold. First, it does not require to train the network before applying pruning, thus greatly reducing training time. Second, it prunes weights while the network is training, making the remaining weights jointly optimize for the task at hand, and the removal of other weights. Third, the design of the schedule helps the training dynamics by pruning very gently at the beginning and also helps the network to recover from pruning by slowly reducing pruning intensity towards the end.


### 4. COMPARISON TO TRADITIONAL SCHEDULES

In this section, we compare our proposed schedule, the One-Cycle Pruning, to other commonly used pruning schedules.

**Pruning Methods.** We compare our pruning technique to several state-of-the-art pruning schedules: One-Shot Pruning, Iterative Pruning and Automated Gradual Pruning (AGP), under a fixed training budget. The optimal training iteration at which the pruning process starts for those schedules, *i.e.* the pretraining phase, is determined by a grid search. We find the optimal starting iteration at 40%, 20% and 20% of training budget for the One-Shot Pruning, Iterative Pruning and AGP, respectively, as depicted in Figure 2.

ResNet-18					VGG-16						
		One-Shot	Iterative	AGP	One-Cycle			One-Shot	Iterative	AGP	One-Cycle
<b>CIFAR-10</b>											
Sparsity	80%	93.10 ± 0.03	93.13 ± 0.03	93.22 ± 0.22	<b>93.49 ± 0.14</b>	90.25 ± 0.14	90.64 ± 0.19	<b>90.87 ± 0.15</b>	90.84 ± 0.09		
	90%	92.42 ± 0.21	91.72 ± 0.08	92.85 ± 0.09	<b>93.31 ± 0.20</b>	89.82 ± 0.19	89.76 ± 0.18	90.67 ± 0.25	<b>90.72 ± 0.40</b>		
	95%	91.58 ± 0.04	87.54 ± 0.39	92.04 ± 0.07	<b>92.76 ± 0.16</b>	89.73 ± 0.37	81.46 ± 2.87	90.56 ± 0.31	<b>90.67 ± 0.11</b>		
<b>CIFAR-100</b>											
Sparsity	80%	74.21 ± 0.09	74.18 ± 0.29	74.78 ± 0.09	<b>74.81 ± 0.16</b>	67.83 ± 0.19	67.80 ± 0.15	67.93 ± 0.06	<b>68.34 ± 0.38</b>		
	90%	73.34 ± 0.23	71.80 ± 0.05	73.83 ± 0.41	<b>74.50 ± 0.24</b>	67.33 ± 0.16	62.66 ± 1.31	67.88 ± 0.39	<b>68.24 ± 0.45</b>		
	95%	71.68 ± 0.16	62.88 ± 0.27	71.92 ± 0.30	<b>73.34 ± 0.21</b>	66.16 ± 0.49	61.95 ± 0.70	<b>67.51 ± 0.19</b>	<b>67.51 ± 0.16</b>		
<b>Caltech-101</b>											
Sparsity	80%	80.31 ± 0.89	79.78 ± 0.56	81.93 ± 0.85	<b>82.31 ± 0.88</b>	77.81 ± 0.96	78.23 ± 0.35	78.45 ± 0.85	<b>78.90 ± 0.88</b>		
	90%	79.87 ± 0.54	77.84 ± 0.31	80.89 ± 0.90	<b>81.84 ± 0.16</b>	78.77 ± 1.06	74.42 ± 2.79	<b>78.57 ± 0.21</b>	78.56 ± 0.31		
	95%	78.57 ± 1.02	73.83 ± 1.28	78.76 ± 1.27	<b>79.81 ± 0.92</b>	76.99 ± 0.78	42.61 ± 2.60	78.68 ± 0.53	<b>78.99 ± 0.50</b>		

**Table 2.** Results of pruning ResNet-18 and VGG-16 with 4 different schedules. Mean and standard deviation of accuracy over 3 rounds are reported. Best results are in bold.



**Fig. 2.** Evolution of sparsity during the training of the 4 studied pruning schedules. Best results are obtained when One-Shot, Iterative and AGP start at 40%, 20% and 20% of training budget, respectively. OCP is applied right from the start.

**Datasets and Architectures.** For our experiments, the datasets have been chosen to be various in terms of image resolution and number of classes. In particular, we evaluate our methods on the three following datasets: CIFAR-10 [19], CIFAR-100 [19] and Caltech-101 [20]. Moreover, those datasets are tested on two types of popular convolutional network architectures: VGG-16 [21] and ResNet-18 [22].

**Training Procedure.** The networks we use for our experiments are trained from a random initialization. The images are first resized to  $224 \times 224$  and are augmented by using horizontal flips, rotations, image warping and random cropping. We train each model using a learning rate warmup until a nominal value of 0.001, then a gradual decay until the end of the training.

**Frameworks and Hardware.** Experiments are conducted using the PyTorch [23] and fastai [24] deep learning libraries for training procedures, fasterai [25] library for the implementation of the pruning methods and using a 12GB Nvidia GeForce GTX 1080 Ti GPU for computation.

**Experiment A.** The first experiment is conducted with a

fixed training budget of 50 epochs. The objective is to identify the pruning schedule that is the most efficient when the training budget is tight. The results, reported in Table 2, show that One-Cycle Pruning usually outperforms other pruning schedules for all the different combinations of studied architectures, datasets and sparsity levels.

**Experiment B.** To better emphasize the impact of the pruning schedule on the training dynamics, we propose to fix a target accuracy and let the training budget change according to the needs of the pruning method in order to reach that accuracy level. For One-Shot and Iterative pruning, the pretraining budget is kept identical, only the fine-tuning budget is extended, *i.e.* the training after the pruning has occurred. We present in Table 3 the results of the training budget required to reach 90%, 70% and 80% of accuracy on CIFAR-10, CIFAR-100, Caltech-101 dataset respectively, using ResNet-18 pruned to a sparsity of 95%. Training budget is expressed relatively to method One-Cycle Pruning.

	One-Shot	Iterative	AGP	One-Cycle
<b>CIFAR-10</b>	2.5×	4×	1×	1×
<b>CIFAR-100</b>	3.33×	7.5×	1.25×	1×
<b>Caltech-101</b>	2×	3.2×	1.4×	1×

**Table 3.** Training budget required to prune ResNet-18 to 95% to achieve a fixed validation accuracy of 90%, 70% and 80% on CIFAR-10, CIFAR-100 and Caltech-101 respectively

**Discussion.** Overall, the technique requiring the most training budget while providing the worst validation accuracy when that budget is fixed is Iterative Pruning. Several papers have also reported a similar phenomenon, indicating that Iterative Pruning required a long fine-tuning time [5, 11]. On the other hand, One-Cycle Pruning seems to outperform other pruning schedules both when the training budget is fixed and

when the target accuracy is fixed, indicating that our proposed schedule is able to reach higher performance faster.

## 5. EXPERIMENTS WITH LOTTERY TICKETS

In order to investigate further in the behaviour of One-Cycle Pruning, we evaluate it using the Lottery Ticket Hypothesis.


**Finding Lottery Tickets** The Lottery Ticket Hypothesis (LTH) states that in each network, there exists a sub-network that, trained in isolation, is able to achieve comparable performance in a comparable training budget as the whole network, as long as they start from the same initialization [1]. To discover such *winning tickets*, authors propose to follow an Iterative Magnitude Pruning (IMP) method that goes as follows: an initial network, designated by a set of weight  $W_0$ , is trained for  $t$  iterations, giving the set of weights  $W_t$ . Then, a binary mask  $m$ , whose values are either 0 if we want to remove the corresponding weight or 1 if we want to keep it, is applied, creating the network  $W_t \odot m$ , where  $\odot$  denotes the element-wise product. The weights are then reinitialized to their original values  $W_0 \odot m$  and the process is performed iteratively, updating  $m$  until the desired sparsity is achieved. The IMP method require several rounds of one-shot pruning to unveil winning tickets. We propose to compare the quality of found tickets when they are uncovered with different pruning schedules.

**Comparison of Tickets.** We conduct the LTH experiments using ResNet-18 and CIFAR-10 with the same hyperparameters as described in Section 4. In the original LTH experiments, IMP was applied in a single step at the end of each training round. In our experiments, we apply the studied schedules during each training round, such that our One-Shot Pruning corresponds to the original IMP experiment. We report the accuracy of the ticket of the last round of pruning in Table 4, *i.e.*  $W_0 \odot m$ , where  $m$  is composed at 95% of 0. From that experiment, we find that untrained sub-networks found with OCP are outperforming the ones found with other schedules. This further confirms that it is possible to find already performant sub-networks directly at initialization and that we do not necessarily need to pretrain a model before starting the pruning process.

	One-Shot	Iterative
Accuracy (%)	$72.213 \pm 0.009$	$85.778 \pm 0.003$
	AGP	OCP
	$86.114 \pm 0.001$	<b><math>89.239 \pm 0.004</math></b>

**Table 4.** Validation accuracy of the reinitialized ResNet-18 sub-networks found by several pruning schedules at their last pruning round, *i.e.* at a sparsity level of 95%. The training is performed on the CIFAR-10 dataset.

**Stability of Tickets.** To further study the difference in the results, we conduct a stability analysis of the found Lottery Tickets [18]. This analysis consists of retraining two copies of the same ticket  $W_0 \odot m$ , leading to two different trained versions  $W_1 \odot m$  and  $W_2 \odot m$ , then linearly interpolate their set of weights, setting a new network with weights  $W_3 \odot m = \alpha(W_1 \odot m) + (1-\alpha)(W_2 \odot m)$  with  $\alpha \in [0, 1]$ . The two copies will then be said linearly connected, *i.e.* they converged to the same linearly connected minimum, if the validation error of  $W_3 \odot m$  remains stable for all the  $\alpha$  values. We report in Figure 3 the evolution of instability against the level of sparsity. We denote instability error the validation error, *i.e.*  $1 - \text{accuracy}$ , with network  $W_3 \odot m$ , whose weights were interpolated at half-way between the two trained copies, *i.e.* for  $\alpha = 0.5$ . We can observe that OCP, as well as AGP, become stable at really low level of sparsity, *i.e.* after performing only a few rounds of LTH, while One-Shot Pruning and Iterative Pruning take more rounds before getting stable. However, AGP seems to be the only one to show signs of increasing instability for high sparsity values, *i.e.* for late LTH pruning rounds.



**Fig. 3.** Evolution of instability error after each round of the Lottery Ticket Hypothesis when using a different pruning schedule. A low instability error indicates that the found ticket  $W_0 \odot m$  is stable to retraining for that particular sparsity level.

## 6. CONCLUSION

In this work, we proposed One-Cycle Pruning, a novel pruning schedule that allows a network to be pruned in a single training-pruning phase, removing the needs of an initial pre-training phase but also of a complex and time-consuming fine-tuning phase. When compared to common pruning schedules, One-Cycle Pruning provides comparable or better results with significantly less computation required. We also studied the quality of Lottery Tickets found with One-Cycle Pruning. We showed that those tickets were of higher quality and were consistently stable both at low and high level of sparsities, further validating the benefits of the proposed pruning schedule.

## 7. REFERENCES

- [1] Jonathan Frankle and Michael Carbin, “The lottery ticket hypothesis: Finding sparse, trainable neural networks,” in *International Conference on Learning Representations, ICLR*, 2019.
- [2] Davis W. Blalock, Jose Javier Gonzalez Ortiz, Jonathan Frankle, and John V. Guttag, “What is the state of neural network pruning?,” in *Machine Learning and Systems, MLSys*, 2020.
- [3] Yann LeCun, John S. Denker, and Sara A. Solla, “Optimal brain damage,” in *Advances in Neural Information Processing Systems, NeurIPS*, 1989.
- [4] Babak Hassibi, David G. Stork, Gregory Wolff, and Takahiro Watanabe, “Optimal brain surgeon: Extensions and performance comparisons,” in *Conference on Neural Information Processing Systems, NeurIPS*, 1993.
- [5] Hao Li, Asim Kadav, Igor Durdanovic, Hanan Samet, and Hans Peter Graf, “Pruning filters for efficient convnets,” in *International Conference on Learning Representations, ICLR*, 2017.
- [6] Y. He, X. Zhang, and J. Sun, “Channel pruning for accelerating very deep neural networks,” in *Conference on Computer Vision, ICCV*, 2017.
- [7] Nathan Hubens and et al., “An experimental study of the impact of pre-training on the pruning of a convolutional neural network,” in *Conference on Applications of Intelligent Systems, APPIS*, 2020.
- [8] Christos Louizos, Max Welling, and Diederik P. Kingma, “Learning sparse neural networks through  $l_0$  regularization,” in *International Conference on Learning Representations, ICLR*, 2018.
- [9] Dmitry Molchanov, Arsenii Ashukha, and Dmitry Vetrov, “Variational dropout sparsifies deep neural networks,” in *International Conference on Machine Learning, ICML*, 2017.
- [10] Trevor Gale, Erich Elsen, and Sara Hooker, “The state of sparsity in deep neural networks,” in *International Conference on Machine Learning, ICML*, 2019.
- [11] Song Han, Jeff Pool, John Tran, and William J. Dally, “Learning both weights and connections for efficient neural networks,” in *Conference on Neural Information Processing Systems, NeurIPS*, 2015.
- [12] Namhoon Lee, Thalaiyasingam Ajanthan, and Philip Torr, “SNIP: Single-shot network pruning based on connection sensitivity,” in *International Conference on Learning Representations, ICLR*, 2019.
- [13] Michael Zhu and Suyog Gupta, “To prune, or not to prune: exploring the efficacy of pruning for model compression,” in *International Conference on Learning Representations, ICLR*, 2018.
- [14] Jonathan Frankle, David J. Schwab, and Ari S. Morcos, “The early phase of neural network training,” in *International Conference on Learning Representations, ICLR*, 2020.
- [15] Alessandro Achille, Matteo Rovere, and Stefano Soatto, “Critical learning periods in deep networks,” in *International Conference on Learning Representations, ICLR*, 2019.
- [16] Aditya Golatkar, Alessandro Achille, and Stefano Soatto, “Time matters in regularizing deep networks: Weight decay and data augmentation affect early learning dynamics, matter little near convergence,” in *Conference on Neural Information Processing Systems, NeurIPS*, 2019.
- [17] Jonathan Frankle, Gintare Karolina Dziugaite, Daniel M. Roy, and Michael Carbin, “The lottery ticket hypothesis at scale,” *abs/1903.01611*, 2019.
- [18] Jonathan Frankle, Gintare Karolina Dziugaite, Daniel Roy, and Michael Carbin, “Linear mode connectivity and the lottery ticket hypothesis,” in *International Conference on Machine Learning, ICML*, 2020.
- [19] Alex Krizhevsky, “Learning multiple layers of features from tiny images,” 2009.
- [20] 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 *Conference on Computer Vision and Pattern Recognition Workshop*, 2004.
- [21] Karen Simonyan and Andrew Zisserman, “Very deep convolutional networks for large-scale image recognition,” in *Conference on Learning Representations, ICLR*, 2015.
- [22] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun, “Deep residual learning for image recognition,” in *Conference on Computer Vision and Pattern Recognition, CVPR*, 2016.
- [23] Adam Paszke and et al., “Pytorch: An imperative style, high-performance deep learning library,” in *Neural Information Processing Systems, NeurIPS*. 2019.
- [24] Gugger S. Howard J., “fastai: A layered api for deep learning,” in *MDPI Information*. 2020.
- [25] Nathan Hubens, “fasterai,” <https://github.com/nathanhubens/fasterai>, 2022.