

Are Straight-Through gradients and Soft-Thresholding all you need for Sparse Training?

Antoine Vanderschueren¹
UCLouvain, Belgium

antoine.vanderschueren@uclouvain.be

Christophe De Vleeschouwer¹
UCLouvain, Belgium

christophe.devleeschouwer@uclouvain.be

Abstract

*Turning the weights to zero when training a neural network helps in reducing the computational complexity at inference. To progressively increase the sparsity ratio in the network without causing sharp weight discontinuities during training, our work combines soft-thresholding and straight-through gradient estimation to update the raw, i.e. non-thresholded, version of zeroed weights. Our method, named **ST-3** for straight-through/soft-thresholding/sparse-training², obtains SoA results, both in terms of accuracy/sparsity and accuracy/FLOPS trade-offs, when progressively increasing the sparsity ratio in a single training cycle. In particular, despite its simplicity, ST-3 favorably compares to the most recent methods, adopting differentiable formulations [42] or bio-inspired neuroregeneration principles [25]. This suggests that the key ingredients for effective sparsification primarily lie in the ability to give the weights the freedom to evolve smoothly across the zero state while progressively increasing the sparsity ratio.*

1. Introduction

State-of-the-art neural networks are composed of a few million parameters, leading to a few billion computations per inference. To limit these computations, sparse networks have been thoroughly investigated in the past few years [43, 4, 27, 41, 21, 38, 3, 6, 32, 22, 42, 25], and significant efforts have been devoted to their efficient hardware implementation [8, 29]. Sparse networks reduce the *inference complexity* by setting a majority of their weight parameters to zero. Straightforward approaches to train sparse networks are based on pruning [43, 9, 4], meaning links are removed from the network (which is equivalent to setting their weight to zero) with no chance of being re-activated

later along the training. Since the pruning decision is generally based on weight magnitudes, this approach penalizes the weights that have to *change signs* during training, since those weights end up being pruned when they cross the zero border [33]. In contrast, modern sparsification methods favor the emergence of zero weights, without definitely canceling the corresponding links, allowing weights to freely switch between an active and inactive state. The wide variety of approaches that have been proposed for this purpose are surveyed in Section 2. Our work proposes a conceptually and computationally simple solution that (surprisingly) quite favorably compares to this rich array of prior art.

Our proposal mitigates computational complexity at training by restricting the learning to a single round of gradient descent iterations. This is in contrast with the iterative post-training pruning solutions [6, 32], which achieve SoA accuracy/sparsity trade-off, but suffer from an extremely large computation cost due to the multiple rounds of training required to progressively increase the pruning ratio.

As many previous works building on a single round of training, our method adopts a dynamic weight thresholding procedure to progressively increase the ratio of weights that are set to zero along the training, and to balance zeroed weights across layers. Our work is however specific in that (i) it continuously updates zeroed weights all along the training, and (ii) it sets the weights to zero based on soft-thresholding, instead of hard-thresholding, popular in previous works. Those specificities are detailed as follows.

To continuously update the zeroed weights, our method takes inspiration from the training of quantized neural networks (QNN), where the raw (unquantized) weights are updated with a Straight-through-estimator (STE) to avoid a gradient that goes to zero due to the step-function of the quantizer [14, 23]. We apply this same principle to the weights that are ‘quantized’ to zero during network sparsification. This means that raw and non-thresholded weights are maintained and updated in the backward path, even if their thresholded version are zeroed in the forward path. STE allows gradients and momentum to switch positive weights towards negative values (and vice-versa), without

¹Part of this work has been funded by the Walloon Region project SmartGate N°1910087, and by the Fonds de la Recherche Scientifique – FNRS

²Source code and weights available at <https://github.com/vanderschuea/stthree>

getting stuck in an inactive state when being close to zero.

To turn low magnitude weights to zero, thereby sparsifying the network, soft-thresholding is adopted. This prevents the abrupt changes of forward-path weight values induced by hard-thresholding operators, which may turn a small gradient update into a sharp weight discontinuity. Such abrupt discontinuities hamper the network accuracy, especially at high sparsity rates, where they lead to premature layer collapse [36], as illustrated in Figure 1.

Despite its simplicity, our approach surpasses the accuracy/sparsity trade-offs achieved by the most recent (and generally more sophisticated) single training cycle alternatives. Thereby, it reveals that the key ingredients required to train effective sparse networks primary consist in (1) giving the weights the capability to evolve freely and smoothly between active and inactive states during training, especially in the early epochs; and (2) progressively increasing the sparsity ratio without causing sharp discontinuities in weight values along the training.

Another convincing experimental argument in favor of our method lies in the fact that, when combined with iterative and thus quite complex (due to multiple rounds of training cycles) post-training pruning solutions proposed in [6, 32], our ST-3 defines a novel SoA accuracy/sparsity trade-off. Overall, due to its effectiveness but also computational and conceptual simplicity, our ST-3 provides a novel and valuable baseline to assess future sparsification methods. Despite the key ingredients of ST-3 having been independently proposed in earlier works (generally to implement multi-cycle pruning or quantization), our work is novel and original in the way it integrates them in a method specifically suited to single-cycle training of sparse networks. Our work is also the first one to demonstrate that those conceptually simple ingredients are sufficient to reach SoA performance, when properly combined.

2. Related Works

While zeroing the weights in a structured manner might lead to higher acceleration when considering specific hardware implementations, the accuracy obtained with unstructured sparsity is generally higher at a given sparsity level [24, 12, 37], and significant speed up can be achieved with unstructured sparsity on off-the-shelf hardware [15, 31, 7, 29]. Hence, this paper concentrates on *unstructured* sparsity. Initial approaches to learn sparse networks simply removed links from the network at the end of an entire training cycle, based on their respective L1 magnitude [19, 9]. This however comes at the cost of a substantial loss in prediction accuracy compared to a dense network, including when the remaining *active* weights undergo a *finetuning* step (i.e. a shorter training cycle, with small learning rate). Modern approaches do not consider the definitive and abrupt removal of links. They instead imple-

ment partial pruning at the end of multiple training cycles or design solutions for progressive weight zeroing.

Reference approaches to progressively increase the zeroed weights ratio have been studied in [43] and [17]. Global Magnitude Pruning (GMP) [43] progressively prunes the links along the training. A related method, named STR, sets the weights that lie below a certain threshold to zero, using soft-thresholding. In contrast to our ST-3, only the non-zeroed weights are updated by gradients.

A few previous methods have revealed the benefit of a *non-permanent* zeroing of weights. RIGL [3] builds on gradient momentum to resurrect some of the zeroed weights. GraNet [25] boosts the pruning plasticity through bio-inspired neuroregeneration principles. ProbMask [42] learns a probabilistic pruning mask that is transformed into a binary one by sampling (multiple times, each requiring gradient computations) a Gumbel distribution, thus making it trainable via gradient descent. Our ST-3 is conceptually and computationally simpler since it relies on straight-through gradient estimation to continuously update zeroed weights. Our experiments also reveal that when combined with soft-thresholding, it reaches better accuracy-sparsity trade-offs than earlier works.

Beyond its specificities regarding weight zeroing (with soft-thresholding) and zeroed weight updates (based on straight-through estimation, as done when training quantized networks [14, 23]), our ST-3 adopts conventional solutions regarding three other important design questions :

Weight significance definition. Most methods use gradient-magnitude [3], weight-magnitude [9, 43, 17], or a mix of both [36] to select the weights to set to zero after some preliminary training. As an alternative, DNW [38] uses neural architecture search to discover neural wirings, i.e. independent critical connections between channels, thereby finding a sparse subnetwork of the complete graph. Our work adopts the main trend, zeroing weights based on their magnitude.

Distributing sparsity across layers. It has been widely documented [30, 3, 20] that the network prediction accuracy is more impacted when pruning occurs close to the input. In a CNN, it is however more beneficial from a computational point of view to prune the first layers since they correspond to high-resolution channels. Two heuristics have been recently proposed to address this trade-off. First, Erdos-Rényi-Kernel (ERK) [30, 3] proposes to scale the global sparsity ratio (= ratio of zeros to the total parameter count) with a layer-wise factor. so as to induce a higher (smaller) sparsity for layers with more (less) parameters, e.g. the first convolution layers of a ResNet-50 remain dense, thereby preserving accuracy. Second, Layer-Aware-Pruning (LAMP) [20] selects the weights to prune globally, but relies on a score that is assigned to each weight rather than on its magnitude. This score is computed on the sorted

and flattened array of weights ($\mathbf{W}_{\text{sorted}}^l$) of each layer (l).

Both ERK and LAMP target higher accuracy at the cost of an increased complexity (as measured in FLOPS). However, our experiments did not reveal significant accuracy gains when using one of those heuristics. Hence, unless said otherwise, weights are considered independently of their layer in our experiment. As detailed in Section 3, to handle use cases where the gain in number of operations (FLOPS) is more important than the accuracy, we propose to normalize the weights by the size of their kernel, which approximates the standard deviation of the weights in the layer kernel. Consequently, this normalization reduces the risk to set a whole kernel to zero, thereby preserving accuracy at high sparsity ratio.

Controlling the global level of sparsity along the training. This third question is the one that most impacts the training computational cost. It is therefore central to our paper, which targets simplicity without penalizing accuracy. Most authors [6, 5, 32, 18] agree on the fact that: (i) the pruning should not start at the beginning of training, due to the *heavy changes* a model undergoes in the early stage of training; and (ii) the increase of sparsity should be progressive, to limit the interference between weight update and weight zeroing. Two strategies co-exist to progressively increase the sparsity ratio.

The first one recommends decoupling training iterations and weight zeroing, to promote training consistency and thus convergence quality. Therefore, it builds on multiple training cycles, and is thus intrinsically more complex. It increments the amount of zero weights by a constant fraction of non-zeroed weights at the end of each full training cycle. The learning rate - and optionally the weights' values - are rewinded (the acronym LRR is used to denote this rewinding process), i.e. reset, to an anterior value, before the training resumes. Doing so, SoA accuracy are obtained, but the training cost largely grows with the sparsity ratio [6, 32]. When adopting the same post-training sparsity increase than LRR, and thus iterating over multiple training cycles, our method outperforms the original LRR method, and defines a novel accuracy/sparsity SoA.

The second strategy primarily targets low computational load, and therefore increases the sparsity ratio along a single training cycle. GMP [43] proposes a cubic increase of the sparsity ratio during training, and has been adopted by ProbMask [42] and GraNet [25]. This increase starts after a set amount of epochs, and finishes before the loss, and thus the gradients, become too small. Among the methods running in a single training cycle, ProbMask [42] and GraNet [25] achieve highest accuracies. As mentioned above, [42], transforms the discrete optimization problem associated to the pruning mask definition into a constrained expected loss minimization problem over a continuous probability space, which is solved using the Gumbel-Softmax trick. This how-

ever comes at the cost of multiple computations of gradients in each training iteration, and our experiments demonstrate that the accuracy in [42] fall behind our approach. GraNet [25], and other similar approaches like [3, 26], update the pruning mask every few thousands iterations, by regenerating connections when pruning others. This plasticity can be seen as a discrete approximation of the regeneration of weights that naturally occurs in our method when adopting soft-thresholding on top of the gradient accumulation inherent to our recommended straight-through estimator.

3. Our Method

This section introduces ST-3 as a new baseline for sparse training. It provides a relevant baseline candidate, because it is easy to implement (see Supplementary File), has very little overhead compared to regular dense training, and performs well with the default dense training parameters. ST-3 adopts weight magnitude soft-thresholding with threshold growing along training to progressively increase the sparsity ratio in the forward path. To preserve the opportunity for a zeroed weight to become significant again, similarly to what is done for quantized network training, it considers a straight-through gradient estimator to update a non-thresholded version of each weight maintained in the backward path. Improved training stability is achieved by scaling the weights in the forward path to compensate for the loss of magnitude inherent to soft-thresholding. Despite ST-3 builds on a set of existing or straightforward solutions (progressive increase of sparsity ratio with magnitude loss compensation, straight-through estimator, and soft-thresholding), it is novel in the sense that no previous work has combined all the ST-3 ingredients in a sparse training context. Our experimental study demonstrates the importance of including all elements to preserve ST-3 performance. They are detailed below, and all contribute to a stable and consistent training, aiming at minimizing the impact of weight zeroing on the gradient descent updates.

Straight-Through-estimation To control the sparsity ratio, our method sets to zero the weights that lie below some adaptive threshold. As mentioned before, a main requirement in sparse training is that the weights that are zeroed at some point during the training, because they are below the threshold, get the opportunity to become significant again in subsequent training iterations. Therefore, our method maintains a dense and continuously updated version of the weights in parallel to their sparse thresholded version, to be considered in the forward path but also to back-propagate the gradients. Following [1, 39], a straight-through estimator is considered to define the gradient with respect to the dense weight as a copy of the gradient with respect to the output of the thresholding operator, namely with respect to the sparse weight. This differentiation of raw (to update)



Figure 1. Accuracy of a sparse ResNet-20 on Cifar-10. Conventional straight-through-estimation (STE), with hard-thresholding, fails at 99% sparsity, as already pointed in [36]. Replacing hard-by soft-thresholding reduces the gradient mismatch and alleviates the issue (Dense accuracy of 91.2% orange dashed line)

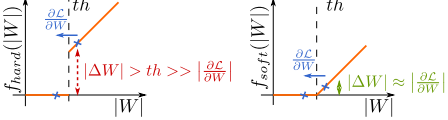


Figure 2. Hard-thresholding (left) induces a discrepancy between the forward weight update and the gradient received in the backward path. Soft-thresholding (right) does not suffer from such inconsistency, since it preserves a smooth relationship between update and gradient value. This prevents sharp weight discontinuities along training.

and thresholded (to use in forward and backward paths) weights is standard in the training of quantized neural networks [14, 28], but is not common in the pruning community.

Soft Thresholding. Most of the persistent pruning literature makes use of hard thresholding to increase the pruning ratio [9, 4, 3]. However, in sparse training, since the weights that are zeroed (because they are below the threshold) at some point during the training might become significant again in subsequent training iterations, they are also likely to cross the pruning threshold, and be subject to the hard thresholding function discontinuity. This might lead to abrupt changes in the forward path, which might be inconsistent with the gradient received in the backward path. Figure 2 depicts a case where the pruning threshold is large compared to the gradient magnitude. In this case, hard thresholding induces severe discontinuities in the forward weight update, while soft thresholding preserves a smooth and consistent evolution of this forward weight. The results shown in Figure 1 confirm that soft thresholding leads to a more stable training, and better accuracy of trained models.

Weight rescaling. The weights that contribute to the same output neuron are referred to as being part of the same *filter*. To preserve the mean absolute magnitude of a neuron in presence of weight zeroing, we propose to rescale the active weights filter-wise, based on the ratio of lost weight-magnitude introduced by the zeroed weights. Formally, letting $W^l \in R^{N_{out} \times N_{in}}$ denote the dense weight matrix (before soft-thresholding) of layer l and th be the sparsity

threshold, then the factor multiplying the soft-thresholded vector associated to the j^{th} output is defined as follows:

$$scale_j^l = \frac{\sum_i^{N_{in}} \|W_{ji}^l\|}{\sum_i^{N_{in}} \begin{cases} \|W_{ji}^l\| & \text{if } \|W_{ji}^l\| > th \\ 0 & \text{otherwise} \end{cases}}$$

This is similar to the rescaling used by dropout [35].

Increase and layer-wise allocation of sparsity. To mitigate the instabilities induced by the thresholding, and leave the network the time to adapt to the deactivation of some weights in the forward path, we increase the sparsity ratio progressively along training. To control the increase of global sparsity, the cubic pattern introduced in [43] is adopted. Two variants of our method have been studied in our experiments to turn the global sparsity ratio into a layer-wise sparsity.

The first variant, denoted **ST-3**, is adopted by default. It simply considers the network weights globally. Hence, no specific attention is drawn to the balance of sparsity across layers. It has the advantage of being straightforward, and our experiments reveal that it results in better accuracy/sparsity tradeoffs as ERK or LAMP.

The second variant, denoted **ST-3^σ**, has been designed to bias the training towards increased FLOPS gains. It arises from the observation that to save FLOPS, it is better to prune weights in layers where the feature maps are larger, since large maps require more operations. However, large feature maps generally manipulate small kernels (i.e. due to the adoption of a small number of high-resolution channels in common architectures), whose initialization results in large weight magnitudes when adopting the popular Kaiming He initialization [11]. As a consequence, as shown experimentally [3], after global thresholding, those kernels with larger weights end-up in being relatively less sparse than the kernels associated to the smaller resolutions channels of wider layers. This aspect is generally overlooked in previous work [6, 4, 32, 20], where only the model compression is considered, disregarding the number of operations affected by the zeroed weights. Hence, most works [17, 42] treat inference acceleration as a by-product, while still optimizing only for a specified sparsity ratio. To favor the zeroing of weights in layers manipulating (few) high resolution channels, and improve the computational gain associated to sparsity, we propose, before soft-thresholding, to multiply each weight by the square root of the number of weights in its corresponding kernel. This can be interpreted as normalizing the weights by their standard deviation in individual kernel since the variance of the weights following the Kaiming He initialization distribution is known to be inversely proportional to the size of the kernel [11].

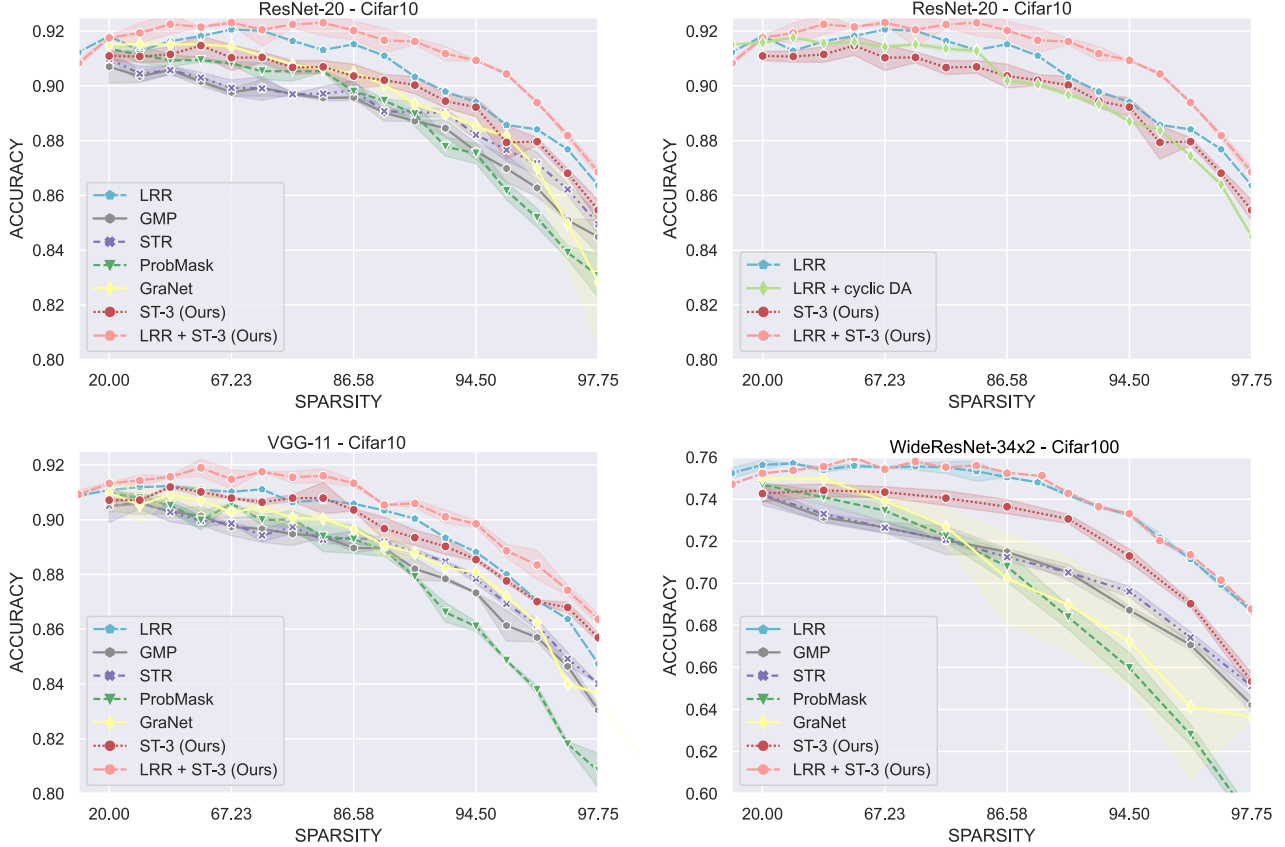


Figure 3. (best viewed in color) (top) Accuracy/sparsity curves obtained, on Cifar10 w/ ResNet-20, (bottom) Accuracy/sparsity curves obtained on Cifar10 w/ VGG-11, and on Cifar100 w/ WideResNet-34x2. The methods starting w/ LRR are recursively trained, i.e. each point on the curve needs every point on its left to be computed first. All points represent the average of 3 runs (with different seeds common to all methods) and standard deviation is shown. The x-axis is depicted in log-scale.

4. Experiments

4.1. Methods used as baselines

Our approach is compared to a set of previous methods that are representative of the SoA and of the recent trends in the field. We now introduce those methods, and explain how they have been configured to provide a fair comparison, and help in quantifying the actual benefit brought by our proposed ST-3.

As a recognized SoA upper bound in terms of accuracy/sparsity trade-off, the Learning-Rate-Rewind (LRR) recursive pruning method is considered [32]. This method however has the obvious downside of requiring significantly more training cycles than every other method considered in this paper. Typically 10 consecutive training cycles are needed to reach 90% sparsity, which isn't practically scalable to larger datasets. In our experiments, LRR is also combined with our ST-3 as follows: the cubic increase of sparsity ratio adopted in ST-3 is replaced by the increase of sparsity ratio implemented by LRR between training cycles,

with learning rate rewind. This hybrid approach is denoted LRR+ST-3. It is as expensive to train than LRR, but leads by far to the best accuracy/sparsity trade-offs, which confirms the benefit of straight-through estimation combined to soft-thresholding with magnitude loss compensation.

Among the methods adopting a single training cycle, the GMP [43], STR [17], and ProbMask [42] methods, presented in Section 2, are probably the most relevant to compare with: GMP and STR because they are conceptually simple threshold-based pruning and sparsification methods, respectively; ProbMask because it corresponds to the current SoA in one-training cycle methods (even if its training time is twice as long as the ST-3 one, due to gradient accumulation required by the Gumbel-Softmax trick at each training step).

For both GMP and STR methods, the *sparsity ratio* increases following the same cubic polynomial function as the one adopted by our method (see Sec. 3), and global thresholding is adopted to achieve the required sparsity ratio, as

for ST-3³. For ProbMask, the sparsity increase ratio and batch-sizes are kept as in the original paper, as changing them led to significant degradation in accuracy.

For completeness, comparison is also provided with DNW [38], RIGL [3] and GraNet [25] (see Section 2 for the presentation of those methods), implemented as recommended by their respective authors.

4.2. Cifar10 and Cifar100

The methods introduced in Section 4.1 have been run on Cifar-10 and Cifar-100 [16]. Three different networks are considered: ResNet-20 [10], VGG-11 [34] (as simplified by [27]), and WideResNet-34x2 [40]. Similar to what is done in related papers, one training cycle consists in 160 epochs and, in the single training cycle scenario, the progressive increase of the sparsity ratio along a training cycle starts at the 5th epoch and ends at the 80th epoch.

One training cycle. In Figure 3, our ST-3 achieves the best accuracy/sparsity tradeoffs among methods trained with a single training cycle, including compared to ProbMask[42]⁴ and GraNet[25]. The soft-thresholding mechanism adopted in STR [17] makes it possible for some weights to move between inactive and active states, which results in a significant boost in performance compared to GMP [43]. However, STR falls short compared to ST-3. This is because STR does not account for gradient information in any way, except for momentum-like terms stored in the optimizer. ST-3 solves this issue, by allowing the accumulation of the gradient information conveyed by the straight-through-estimator.

Multiple training cycles. The computational cost inherent to the multiple training cycles imposed by LRR makes it impractical. However, LRR remains a de-facto upper bound to compare with when considering the training of sparse network. Without surprise, we observe in Figure 3 that LRR performs better than ST-3. More interestingly, LRR+ST-3 method is even better, defining a new SoA. The superiority of LRR+ST-3 over LRR provides an additional clue that the use of a straight-through estimator to continuously and smoothly update the raw (non-thresholded) version of zeroed weights leads to a significant increase in performance when training sparse networks. A deeper comparison between ST-3 and LRR reveals that ST-3 gets closer to LRR on

Cifar10 than with Cifar100. We attribute this phenomenon to the fact that the increased data augmentation associated to the longer training implemented by LRR is more beneficial for Cifar100 (only 500 training images per class) than for Cifar10 (5000 training images per class). To challenge this hypothesis, we have trained ST-3 for 4x more epochs than a reference cycle at a sparsity threshold of 67.23% on ResNet-20. This longer training has resulted in a test accuracy of 92.38% which is better than the 92.22% obtained by LRR with 6x the number of epochs of a reference training cycle. The role of data augmentation in the performance improvement observed with longer training is also explored on ResNet-20 in Figure 3 (*top right*). LRR is applied as described earlier, but the augmentation seed is reset between each training cycle, ensuring that LRR manipulates the same augmented samples than one-cycle methods. With this additional constraint, we observe that the LRR performance degrades up to the one of our ST-3.

4.3. ImageNet

ResNet-50 [10] and MobileNetv1 [13] have been trained on ImageNet [2] using the same standard hyper-parameters, number of epochs (=100), and data-augmentation as previous related works [3, 17, 42, 25], leading to a *dense* baseline accuracy of 77.1%. To train sparse models, only one-cycle methods have been considered, due to the excessive training time required by the LRR iterations. The sparsity progressively increases between the 5th and the 50th epochs. Values for curves corresponding to related works with ResNet-50 on ImageNet are directly copied from their respective papers⁵.

Accuracy/Sparsity In Figure 4, ST-3 reaches the best accuracy/sparsity tradeoffs. The closest contenders are the most recent results provided by GraNet and ProbMask, but a consistent gap remains that widens up to 5% at 95% sparsity. Our methods unequivocally pull ahead when optimizing sparsity only. It is interesting to note that the gap between the two variants of our method, namely ST-3 and ST-3^σ, is moderate and stays more or less constant across all sparsity ratios. The superiority of ST-3 over ST-3^σ confirms previous observations indicating that pruning earlier layers penalizes accuracy (slightly) more than zeroing later ones [3].

Accuracy/FLOPS FLOPS is a widespread metric to assess the computational gain corresponding to an ideal scenario where every multiplication with zero can be avoided.

³Homogenizing the methods configuration does not significantly impact the method performance (as demonstrated by our results on ImageNet in Sec. 4.3)

⁴Note that, in contrast to what is implemented in the code released with the original paper, our results have been generated by adapting the network’s batch-normalization moving-mean and moving-variance on the training – and not the testing – set. Adapting on the testing set indeed raises a methodological issue. This difference explains the slight discrepancy between our plots and the ones in [42].

⁵In particular, this means that the ProbMask accuracy is the one obtained in [42], i.e. when using test samples to define the batch-normalization parameters. Hence, they correspond to an upper bound of the results that would be obtained when defining batch-normalization parameters with training samples only.

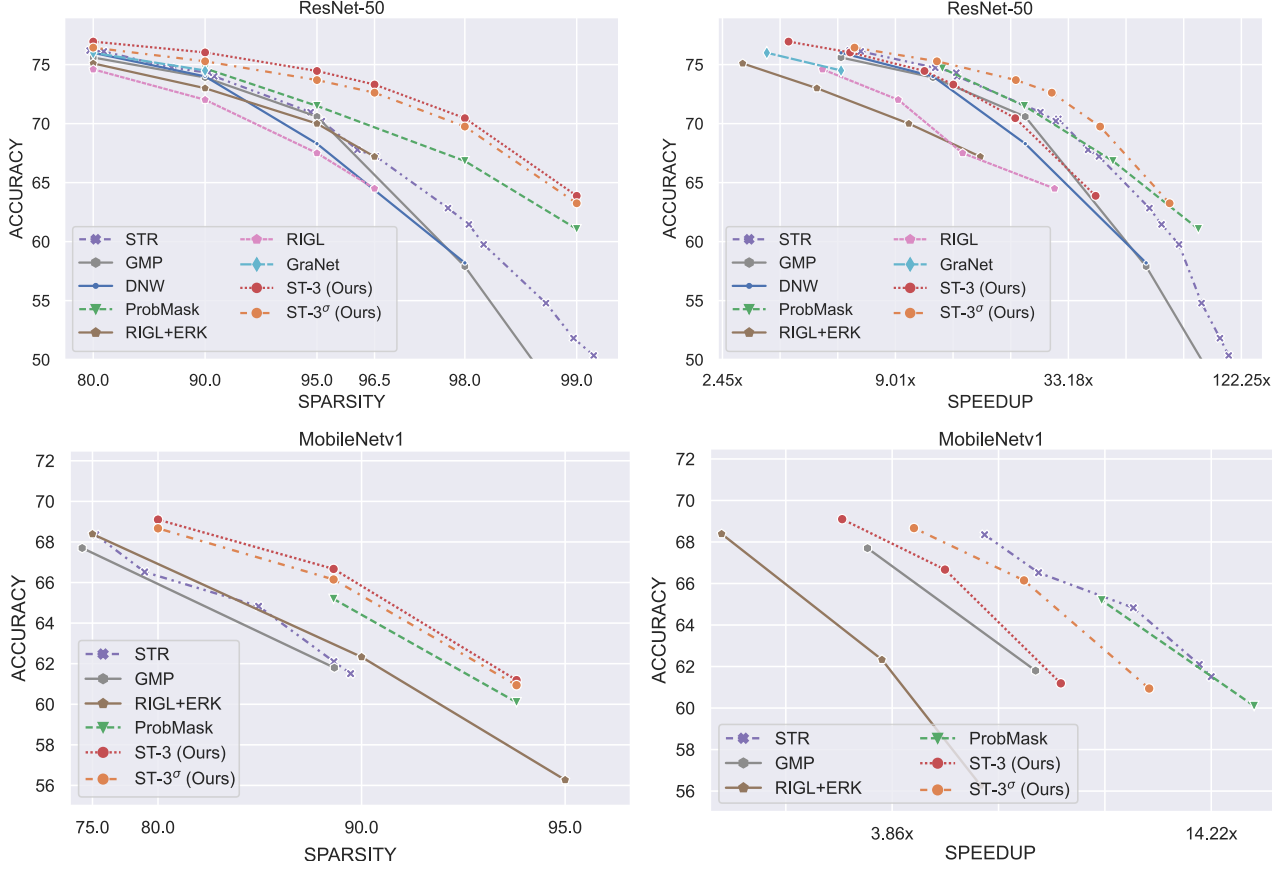


Figure 4. Accuracy on ImageNet w/ ResNet-50 (top) and MobileNetV1 (bottom) as a function of the global sparsity (left), and of the FLOPS (right). All graphs use a log-scale for the x-axis.

In practice, inference speedup is measured as the FLOPS reduction factor. As depicted in Figure 4, our ST-3^σ method, a ST-3 variant proposed to bias the distribution of sparsity across layers in favor of layers with higher resolution, significantly improves sparsity/speedup tradeoffs, making our approach competitive in that respect as well.

4.4. Ablation studies

Isolating ST-3 components Figure 5 evaluates the importance of the three components of our method, namely straight-through-estimation (STE), soft-thresholding, and weight rescaling for mean magnitude loss compensation. Straight-through-estimation and soft-thresholding appear to be the main contributors to the success of our method. The rescaling factor helps in some respect but only when soft-thresholding is used. When used with pure STE, meaning with hard-thresholding, it tends to amplify the number of sharp discontinuities in the evolution of weights, thereby causing a severe penalty.

Weight switching between active/inactive states Figure 6 investigates how the soft-thresholded weights in ST-3

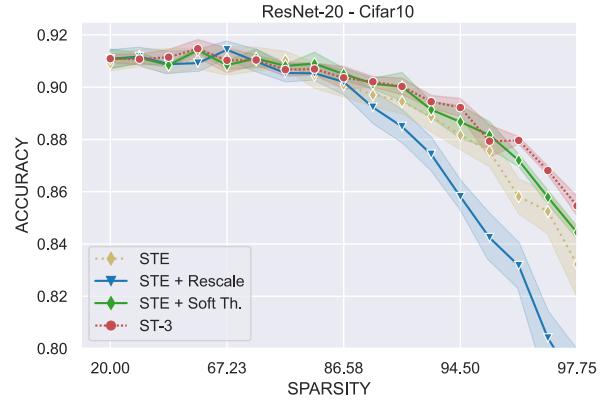


Figure 5. A study dissecting the impact of straight-through estimation, soft-thresholding and weight-rescaling during sparse training. A log-scale is used for the x-axis.

switch between zero and non-zero values during training. It presents how the counts of switches per weight are distributed in four sequences of epochs, obtained by partitioning the total number of epochs in groups of 40 epochs. For clarity, weights that switch less than two times are ignored

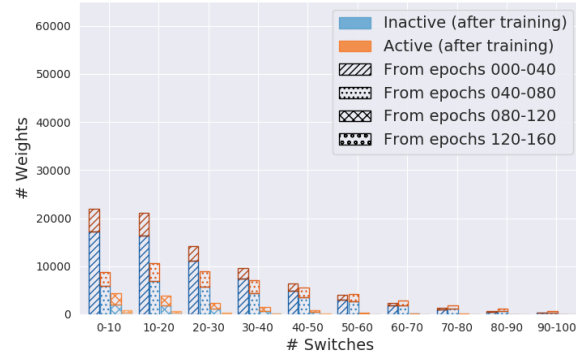
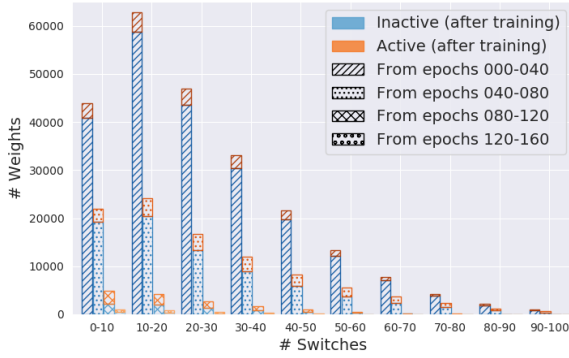


Figure 6. (best viewed in color) Number of switches between active and inactive state for the weights of a ResNet-20 trained with ST-3 for 160 epochs on Cifar-10 at 90% sparsity. The global sparsity ratio is fixed (left) or cubically increasing (right). (Weights that switch once or never are not included)

in the plot. When plotting these distributions, a color code is used to distinguish the weights that are active (i.e. non-zero) at the end of training from the ones that are zeroed.

On the left, the sparsity ratio is set to its final value from the beginning of training, while it progressively increases along training on the right. We observe that progressively increasing the sparsity ratio fundamentally changes the weight switching pattern. Progressive increase stabilize the training, as attested by the reduced amount of switches between states, especially in the second group of epochs. It also makes the state switching pattern closer to the typical sign switching pattern observed in dense training (the signs of weights that remain significant at the end of training change more often in the first epochs). A progressive increase of the sparsity is thus desired since it complements the stability offered by straight-through estimation and soft-thresholding.

5. Conclusion

Throughout this paper the relevance of our straightforward and computationally simple approach, ST-3^(σ), has been widely demonstrated. Its efficiency in training very sparse networks has been assessed on both cifar-10(0) and ImageNet with a variety of architectures, attaining single-cycle SoA accuracy. The fact that our simple and relatively straightforward approach achieves the highest accuracies among approaches with comparable complexity suggests that the key towards effective sparsification does not hold in a sophisticated formulation of the optimization problem but rather in the ability to give the weights the freedom to evolve across the zero state while progressively increasing the sparsity ratio.

References

- [1] Yoshua Bengio, Nicholas Léonard, and Aaron C. Courville. Estimating or propagating gradients through stochastic neurons for conditional computation. *CoRR*, abs/1308.3432, 2013.
- [2] 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*, pages 248–255. Ieee, 2009.
- [3] Utku Evci, Trevor Gale, Jacob Menick, Pablo Samuel Castro, and Erich Elsen. Rigging the lottery: Making all tickets winners. *37th International Conference on Machine Learning, ICML 2020*, PartF16814:2923–2933, 2020.
- [4] Jonathan Frankle and Michael Carbin. The lottery ticket hypothesis: Finding sparse, trainable neural networks. *7th International Conference on Learning Representations, ICLR 2019*, pages 1–42, 2019.
- [5] Jonathan Frankle, Gintare Karolina Dziugaite, Daniel Roy, and Michael Carbin. Pruning neural networks at initialization: Why are we missing the mark? *ICLR*, 2021.
- [6] Jonathan Frankle, David J Schwab, and Ari S Morcos. The Early Phase of Neural Network Training. In *International Conference on Learning Representations*, pages 1–20, 2020.
- [7] Trevor Gale, Erich Elsen, and Sara Hooker. The state of sparsity in deep neural networks. *CoRR*, abs/1902.09574, 2019.
- [8] Scott Gray, Alec Radford, and Diederik P Kingma. GPU Kernels for Block-Sparse Weights. *OpenAI '17*, (1):12.
- [9] Song Han, Jeff Pool, John Tran, and William J. Dally. Learning both weights and connections for efficient neural networks. *Advances in Neural Information Processing Systems*, 2015-Janua:1135–1143, 2015.
- [10] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep Residual Learning for Image Recognition. *37(50):1951–1954*, dec 2015.
- [11] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level perfor-

- mance on imagenet classification. *CoRR*, abs/1502.01852, 2015.
- [12] Yang He, Ping Liu, Ziwei Wang, Zhilan Hu, and Yi Yang. Filter pruning via geometric median for deep convolutional neural networks acceleration. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2019-June:4335–4344, 2019.
 - [13] 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. *arXiv preprint arXiv:1704.04861*, 2017.
 - [14] Itay Hubara, Matthieu Courbariaux, Daniel Soudry, Ran El-Yaniv, and Yoshua Bengio. Quantized Neural Networks: Training Neural Networks with Low Precision Weights and Activations. 2016.
 - [15] Nal Kalchbrenner, Erich Elsen, Karen Simonyan, Seb Noury, Norman Casagrande, Edward Lockhart, Florian Stimber, Aäron Van Den Oord, Sander Dieleman, and Koray Kavukcuoglu. Efficient neural audio synthesis. *35th International Conference on Machine Learning, ICML 2018*, 6:3775–3784, 2018.
 - [16] Alex Krizhevsky. (CIFAR10) Learning Multiple Layers of Features from Tiny Images. ... *Science Department, University of Toronto, Tech. ...*, pages 1–60.
 - [17] Aditya Kusupati, Vivek Ramanujan, Raghav Somani, Mitchell Wortsman, Prateek Jain, Sham Kakade, and Ali Farhadi. Soft threshold weight reparameterization for learnable sparsity. pages 5544–5555, 2020.
 - [18] Guillaume Leclerc and Aleksander Madry. The Two Regimes of Deep Network Training. 2020.
 - [19] Yann LeCun, John S Denker, and Sara A. Solla. Optimal Brain Damage (Pruning). *Advances in neural information processing systems*, pages 598–605, 1990.
 - [20] Jaeho Lee, Sejun Park, Sangwoo Mo, Sungsoo Ahn, and Jinwoo Shin. Layer sparsity for the magnitude-based pruning. *CoRR*, abs/2010.07611, 2020.
 - [21] Namhoon Lee, Thalaiyasingam Ajanthan, and Philip H.S. Torr. SnIP: Single-shot network pruning based on connection sensitivity. In *7th International Conference on Learning Representations, ICLR 2019*, 2019.
 - [22] Bailin Li, Bowen Wu, Jiang Su, and Guangrun Wang. EagleEye: Fast Sub-net Evaluation for Efficient Neural Network Pruning. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 12347 LNCS:639–654, 2020.
 - [23] Fengfu Li, Bo Zhang, and Bin Liu. Ternary Weight Networks. (Nips), 2016.
 - [24] Hao Li, Hanan Samet, Asim Kadav, Igor Durdanovic, and Hans Peter Graf. Pruning filters for efficient convnets. *5th International Conference on Learning Representations, ICLR 2017 - Conference Track Proceedings*, (2016):1–13, 2017.
 - [25] Shiwei Liu, Tianlong Chen, Xiaohan Chen, Zahra Atashgahi, Lu Yin, Huanyu Kou, Li Shen, Mykola Pechenizkiy, Zhangyang Wang, and Decebal Constantin Mocanu. Sparse Training via Boosting Pruning Plasticity with Neuroregeneration. (NeurIPS), 2021.
 - [26] Shiwei Liu, Lu Yin, Decebal Constantin Mocanu, and Mykola Pechenizkiy. Do we actually need dense over-parameterization? in-time over-parameterization in sparse training. In *International Conference on Machine Learning*, pages 6989–7000. PMLR, 2021.
 - [27] Zhuang Liu, Mingjie Sun, Tinghui Zhou, Gao Huang, and Trevor Darrell. Rethinking the value of network pruning. *7th International Conference on Learning Representations, ICLR 2019*, pages 1–21, 2019.
 - [28] Naveen Mellempudi, Abhisek Kundu, Dheevatsa Mudigere, Dipankar Das, Bharat Kaul, and Pradeep Dubey. Ternary Neural Networks with Fine-Grained Quantization. 2017.
 - [29] Asit K. Mishra, Jorge Albericio Latorre, Jeff Pool, Darko Stosic, Dusan Stosic, Ganesh Venkatesh, Chong Yu, and Paulius Micikevicius. Accelerating sparse deep neural networks. *CoRR*, abs/2104.08378, 2021.
 - [30] Decebal Constantin Mocanu, Elena Mocanu, Peter Stone, Phuong H. Nguyen, Madeleine Gibescu, and Antonio Liotta. Scalable training of artificial neural networks with adaptive sparse connectivity inspired by network science. *Nature Communications*, 9(1):1–12, 2018.
 - [31] Jongsoo Park, Hai Li, Sheng Li, Wei Wen, Yiran Chen, Ping Tak Peter Tang, and Pradeep Dubey. Faster cnns with direct sparse convolutions and guided pruning. *5th International Conference on Learning Representations, ICLR 2017 - Conference Track Proceedings*, 2017:1–12, 2017.
 - [32] Alex Renda, Jonathan Frankle, and Michael Carbin. Comparing Rewinding and Fine-tuning in Neural Network Pruning. (2019):1–31, 2020.
 - [33] Victor Sanh, Thomas Wolf, and Alexander M. Rush. Movement pruning: Adaptive sparsity by fine-tuning. *Advances in Neural Information Processing Systems*, 2020-Decem(NeurIPS):1–12, 2020.
 - [34] K. Simonyan and A. Zisserman. Very Deep Convolutional Networks for Large-Scale Image Recognition. *CoRR*, 2014.
 - [35] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: A simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15(56):1929–1958, 2014.
 - [36] Hidenori Tanaka, Daniel Kunin, Daniel L.K. Yamins, and Surya Ganguli. Pruning neural networks without any data by iteratively conserving synaptic flow. *Advances in Neural Information Processing Systems*, 2020-Decem(NeurIPS), 2020.
 - [37] Huan Wang, Can Qin, Yulun Zhang, and Yun Fu. Neural Pruning via Growing Regularization. pages 1–15, 2020.
 - [38] Mitchell Wortsman, Ali Farhadi, and Mohammad Rastegari. Discovering neural wirings. *Advances in Neural Information Processing Systems*, 32(NeurIPS), 2019.
 - [39] Penghang Yin, Jiancheng Lyu, Shuai Zhang, Stanley J. Osher, Yingyong Qi, and Jack Xin. Understanding straight-through estimator in training activation quantized neural nets. *ICLR*, 2019.
 - [40] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. *CoRR*, abs/1605.07146, 2016.
 - [41] Hattie Zhou, Janice Lan, Rosanne Liu, and Jason Yosinski. Deconstructing Lottery Tickets: Zeros, Signs, and the Supermask. (NeurIPS), 2019.

- [42] Xiao Zhou, Weizhong Zhang, Hang Xu, and Tong Zhang. Effective Sparsification of Neural Networks with Global Sparsity Constraint. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 3598–3607, 2021.
- [43] Michael H. Zhu and Suyog Gupta. To prune, or not to prune: Exploring the efficacy of pruning for model compression. *6th International Conference on Learning Representations, ICLR 2018 - Workshop Track Proceedings*, 2018.