

---

# ESPN: EXTREMELY SPARSE PRUNED NETWORKS

---

Minsu Cho    Ameya Joshi    Chinmay Hegde\*  
 Tandon School of Engineering  
 New York University  
 {mc8065, aaj458, chinmay.h}@nyu.edu

## ABSTRACT

Deep neural networks are often highly overparameterized, prohibiting their use in compute-limited systems. However, a line of recent works has shown that the size of deep networks can be considerably reduced by identifying a subset of neuron indicators (or mask) that correspond to significant weights prior to training. We demonstrate that a simple iterative mask discovery method can achieve state-of-the-art compression of very deep networks. Our algorithm represents a hybrid approach between single shot network pruning methods (such as SNIP) with Lottery-Ticket type approaches. We validate our approach on several datasets and outperform several existing pruning approaches in both test accuracy and compression ratio.

## 1 Introduction

**Motivation.** Neural networks have achieved state of the art results across several domains such as computer vision, language processing, and reinforcement learning. This performance is generally contingent on large, over-parameterized networks that are trained using massive amounts of data. For example, the current state of the art on the ImageNet classification task uses a network with over 480 million parameters [TVDJ20], and consequently, the best performing networks are prohibitive in terms of computational and memory requirements. Therefore, compressing neural networks is vital for resource limited settings (such as self-driving cars and mobile devices). In this work, we present an algorithm for compressing neural networks to far higher degree of sparsity levels than has been reported in the literature.

**Challenges.** Network pruning involves finding a smaller, more efficient representation of a given reference neural network. Pruning strategies include quantization [HMD15, HCS<sup>+</sup>17], sparsification [Cha89, TLFF18, DZW18, FC18, LAT19], and distillation [BFG18, SB16, PPA18]. See [CWZZ17] for a detailed look at several approaches towards neural network compression.

In sparsification-based network pruning, the goal is to select a subset of the original weights that are easily encoded while preserving nominal performance on a test set. However, such methods involve expensive *retraining* of the weights once the subset is identified, as well as requires significant hyper-parameter tuning. [LAT19] propose a method called SNIP, which identifies a salient subset of weights for a given dataset *prior* to training. This avoids re-training of the weights and thereby offers computational advantage over existing methods, but achieves weaker compression performance as a trade-off.

The authors of [FC18] present the *Lottery Ticket hypothesis*, wherein they demonstrate the existence of randomly initialized sub-networks whose weights can be optimized to achieve high performance. Subsequent work by [FDRC19] stabilizes the algorithm for finding a winning ‘ticket’ (i.e., a good sub-network) via iterative pruning methods. However, iterative methods incur larger computation compare to single-shot methods such as [LAT19]. [ZLLY19] further suggest that finding the winning ticket for very deep networks is an implicit consequence of the training mechanism. In any case, all these methods require higher computational costs, as well as significant hyperparameter tuning overhead, than those incurred by a baseline training approach.

---

\*This work was supported in part by grants CCF-2005804 and CCF-1815101 from the National Science Foundation.

**Contributions.** Given these observations, clear trade-offs exist among computation, sensitivity to hyperparameters, and final test performance. In this paper, we show that it is possible to achieve excellent pruning results with not only low computational costs, but also while being robust to tuning hyperparameters. In essence, our method merges [LAT19]’s single-shot pruning approach with the iterative pruning strategy by [ZLLY19]. We demonstrate that our approach is successful at pruning well-initialized, large networks such as residual networks (ResNets). Somewhat surprisingly, our empirical results demonstrate that it may be possible to compress neural networks to considerably higher pruning levels than has been reported so far in the literature.

Our specific contributions are fourfold:

1. We present an algorithm for training extremely sparse neural networks, by learning masking operators for the weights using gradient updates through convex relaxation.
2. We achieve a 5X improvement over baseline methods such as SNIP, while conserving accuracy.
3. We provide extensive comparison on our approach on a variety of models and datasets, outperforming various state-of-the-art sparse pruning methods especially in extreme pruning ratio ( $> 99\%$ ).
4. Finally, we conduct ablation studies and hyperparameter sensitivity experiments to analyse each component of our approach in detail.

## 2 Related Work

Early neural network compression approaches [HMD15, CK14, GYC16] rely on simple magnitude-based heuristic pruning tactics and show significant compression of deep networks while preserving performance. [HMD15], [GYC16], and [ZG17] rely on iterative fine-tuning and pruning to reduce the effect of removing weights on the output. Other work adopts second-order approximation of loss surface ([LDS90], [HS93]) to minimize performance degradation, but do not scale well to larger models. Leveraging techniques from information theory and Bayesian statistics have further improved model compression ratios, including information bottlenecks [DZW18], Fisher pruning [MTK<sup>+</sup>17, TKTH18], entropy-penalized reparameterization [OBSS20], and Variational Bayes [MAV17, LUW17, UMW17].

While most of the above approaches are applied to pre-trained existing networks, there has been a recent surge in approaches that interleave pruning and training. Such approaches rely on reparameterizing the weights and training over the modified network. Methods such as low-rank decomposition [DSD<sup>+</sup>13], FastFood transforms [YMD<sup>+</sup>15], hashing [CWT<sup>+</sup>15] leverage reparameterization schemes for fully connected layers. An obvious reparameterization is inducing sparsity on model connections. However, training sparse models can be unstable. To resolve such issues, [MMS<sup>+</sup>18] and [DZW18] use iterative parameter growth, allowing the sparse architecture additional degrees of freedom. Evolutionary approaches such as [BKML18] propose a dynamic sparsification model using connectivity constraints to train an implicitly sparse model. [MW19] further improve this by allowing a global sparsity constraint instead of a layer-wise approach, thus forcing a network to learn a consistent sparse model.

Another branch of pruning algorithms learns the auxiliary (masking) parameters via gradient-based approaches by relaxing non-differentiable constraints (e.g.,  $\ell_0$ -constraints) to differentiable estimators. [LWK18] re-parameterizes the weights with the element-wise product of its weight and auxiliary parameters sampled by learning the hard concrete distribution. [XWR19] adopts the straight-through estimator to take gradients with respect to the indicator function which generates the binary masks to train auxiliary parameters. [ZLLY19] and [SSM19] trains the auxiliary parameters through gradient descent by reparameterizing with Bernoulli samplers with sigmoidal probabilities on the auxiliary parameters. [SWR20] adopts the hard-concrete distribution proposed in [LWK18] to learn both weights and auxiliary parameters simultaneously. We note that our work and the recent preprint [SWR20] share similar concepts of updating weights and masks simultaneously, but take different approaches for treating non-differentiable regularized terms.

[LRLZ17] present a runtime pruning method that models the flow of activations through the network as a Markov process, and use reinforcement learning to learn the best policy for pruning. [HLL<sup>+</sup>18] also use reinforcement learning in the context of AutoML to train model architectures that are explicitly optimized for edge devices.

Instead of pruning in the middle of or after the training, [LAT19, WZG20, VSF20] proposes the method pruning network prior to the training based on saliency scores from an untrained network. While pruning an untrained network provides a huge advantage in computational cost, the pruning algorithms on a trained network provides par or better compressing performances than single-shot prunings on untrained networks.

While pruning weights in the level of each weight (unstructured) has the advantage of compressing the network in the highest degree of freedom, the limitation on inference speed exist due to limited support in current deep learning software frameworks on parallel processing in each weight level. Instead, structured pruning approaches [LLS<sup>+</sup>17, LSZ<sup>+</sup>19, LWL17, GRK17, AHS17] focuses pruning in a higher structured level such as neurons or filters such that

corresponding weight matrix only contains zero elements. In this work, we focus on unstructured weight pruning by targeting competitive sparsity-accuracy tradeoff.

### 3 Preliminaries

**Notation.** Let  $f : \mathbb{R}^d \rightarrow \mathbb{R}^k$  be a neural network with weights  $\mathbf{w} \in \mathbb{R}^m$ . Assume  $c_i \in \{0, 1\}$ , an auxiliary indicator for each of the weight, where each  $c_i$  indicates if the weight  $w_i$  is a *salient* parameter. Saliency here refers to the effect of zeroing out  $w_i$  on the final risk. A high saliency value (greater than some threshold  $\epsilon$ ) will be indicated by  $c_i = 1$ .

To motivate our proposed approach, we first describe two existing pruning methods.

**SNIP.** Single-Shot Network Pruning (SNIP) [LAT19] introduces a saliency based criterion to prune network connections. The premise is that since neural networks are often overparameterized, they have redundant connections which can be identified prior to training using a saliency-based approach.

Given a dataset  $\mathcal{D} = \{\mathbf{x}, \mathbf{y}\}$ , and a desired network sparsity level  $k$ , training a sparse neural network amounts to solving the following optimization problem:

$$\min_{\mathbf{w}} L(f(\mathbf{w}; \mathbf{x}), \mathbf{y}), \quad \text{s.t. } \|\mathbf{w}\|_0 \leq k. \quad (3.1)$$

Here,  $L(\cdot)$  refers to any standard loss function, such as the cross entropy or  $\ell_2$  loss. Enforcing the  $\ell_0$ -constraint is combinatorially difficult but can be relaxed either via a sparsifying penalty [CPI18, Cha89, Set97, LWK18] or a saliency-based weight dropping mechanism [HMD15, MS89].

SNIP instead uses the following alternative formulation using an auxiliary indicator variable,  $\mathbf{c} \in \{0, 1\}^m$  that indicates if a weight contributes (positively or negatively) to the output. The modified loss therefore is as follows,

$$\min_{\mathbf{w}, \mathbf{c}} L(f(\mathbf{c} \odot \mathbf{w}; \mathbf{x}), \mathbf{y}), \quad \text{s.t. } \|\mathbf{c}\|_0 \leq k. \quad (3.2)$$

Here,  $\odot$  refers to element-wise multiplication. Observe that the auxiliary variable decouples the constraint from the weights,  $\mathbf{w}$ . This decoupling allows for measuring the saliency of a specific weight matrix by measuring the effect of removing a connection. For example, if we consider the difference in loss by removing the  $j^{th}$  edge:

$$\Delta L_j = L(\mathbf{1} \odot \mathbf{w}) - L((\mathbf{1} - \mathbf{e}_j) \odot \mathbf{w}) \quad (3.3)$$

where  $\mathbf{e}_j$  is an indicator vector, then the absolute value of  $\Delta L_j$  indicates the contribution of the  $j^{th}$  parameter to the performance of the network. Instead of directly computing  $\Delta L$  for  $m$  weights which involves cumbersome forward passes, the author approximates  $\Delta L$  with the gradient,  $\partial L / \partial c_j$ , which can be easily calculated using automatic differentiation. *A posteriori*, only the top  $k$  connections with the highest gradient values are retained. This can be done using a simple thresholding operation. Specifically,  $c_j$  is defined as the following indicator variable (denoting connection sensitivity),

$$c_j = \mathbb{1}[s_j - s_k \geq 0], \quad \text{where } s_j = \frac{\partial L / \partial c_j}{\sum_i \partial L / \partial c_i}, \quad (3.4)$$

and  $s_k$  corresponds to the  $k^{th}$  largest sensitivity value. The network weights are then randomly initialized and further trained with the  $c_j$ 's kept constant. This allows for a single-shot compression scheme with only one round of training.

**Lottery Ticket Hypothesis.** The Lottery Ticket (LT) hypothesis proposed by [FC18] claims the existence of a smaller subnetwork within a standard large, dense neural network architecture that will provide competitive performance when trained from scratch as that of the original network. This suggests the need for network pruning to become an essential component of the training process. However, in order to find such a subnetwork, the authors propose an iterative process that alternately prunes and retrains the pruned network from scratch. They also introduce ‘rewinding’, where one retrains the pruned model starting with the initial (random) weights instead of the final learned weights upon which the pruning is performed.

[ZLLY19] further systematically analyse the LT hypothesis and provide two important observations: (1) new values on kept weights should have the same sign as the original initial values (in line with why rewinding in [FC18] is important), (2) it is important to set to masked weights to zero, instead of using values other than zero. From the second observation, [ZLLY19] suggests that magnitude-based masking criteria (which [FC18] use) tends to prune weights that seem to move towards zero during training.

---

**Algorithm 1** ESPN-FINETUNE

---

```
1: Inputs:  $f(\mathbf{w})$ : pre-trained network ( $\mathbf{w} \in \mathbb{R}^d$ ),  $\mathbf{c} = \mathbf{1}$ : auxiliary parameters,  $T$ : fine-tuning epochs,  $\alpha$ : sparsity weight,  $\eta$ : learning rate,  $p$ : pruning ratio,  $\epsilon$ : threshold.
2: procedure TRAINING  $w$  AND  $c$  VIA SGD
3:    $\bar{\mathbf{w}} \leftarrow (\mathbf{w}, \mathbf{c})$ 
4:   while  $N_c \geq d \cdot (1 - p)$  do
5:      $\bar{\mathbf{w}} \leftarrow \bar{\mathbf{w}} - \eta \nabla_{\bar{\mathbf{w}}} (L(f(\mathbf{w} \odot \mathbf{c}; x), y) + \alpha \|\mathbf{c}\|_1)$ 
6:      $N_c \leftarrow \text{SUM}(\mathbb{1}(\mathbf{c} > \epsilon))$ 
7:   end while
8:    $\mathbf{w} \leftarrow \mathbf{w} \odot \mathbf{c}$ 
9:    $\mathbf{c} \leftarrow \mathbb{1}(\mathbf{c} > \epsilon)$ 
10: end procedure
11: procedure FINE-TUNE THE NETWORK
12:   Train  $f(\mathbf{w} \odot \mathbf{c})$  respect to  $\mathbf{w}$  for  $T$  epochs.
13: end procedure
```

---

---

**Algorithm 2** ESPN-REWIND

---

```
1: Inputs:  $f(\mathbf{w})$ : Untrained Network ( $\mathbf{w} \in \mathbb{R}^d$ ),  $t$ : warmup epochs,  $T$ : epochs  $\mathbf{c} = \mathbf{1}$ : auxiliary parameters,  $\alpha$ : Lasso coefficient,  $\eta$ : learning rate,  $p$ : pruning ratio,  $\epsilon$ : threshold.
2: procedure WARMUP TRAINING
3:   for epoch  $\in \{1, \dots, t\}$  do
4:      $\mathbf{w} \leftarrow \mathbf{w} - \eta \nabla_{\mathbf{w}} (L(f(\mathbf{w})))$ 
5:   end for
6:    $\mathbf{w}_t \leftarrow \mathbf{w}$ 
7: end procedure
8: procedure TRAINING  $\mathbf{w}$  AND  $\mathbf{c}$  VIA SGD
9:    $\bar{\mathbf{w}} \leftarrow (\mathbf{w}, \mathbf{c})$ 
10:  while  $N_c \geq d \cdot (1 - p)$  do
11:     $\bar{\mathbf{w}} \leftarrow \bar{\mathbf{w}} - \eta \nabla_{\bar{\mathbf{w}}} (L(f(\mathbf{w} \odot \mathbf{c}; x), y) + \alpha \|\mathbf{c}\|_1)$ 
12:     $N_c \leftarrow \text{SUM}(\mathbb{1}(\mathbf{c} > \epsilon))$ 
13:  end while
14:   $\mathbf{c} \leftarrow \mathbb{1}(\mathbf{c} > \epsilon)$ 
15: end procedure
16: procedure REWIND AND TRAIN THE NETWORK
17:   $\mathbf{w} \leftarrow \mathbf{w}_t \odot \mathbf{c}$ 
18:  Train  $f(\mathbf{w})$  respect to  $\mathbf{w}$  via SGD for  $T - t$  epochs
19: end procedure
```

---

## 4 Our Approach: ESPN

While SNIP does enable very good network compression at moderate computational cost, a single-shot approach before training has several disadvantages. The primary issue is that connection sensitivities estimated using single-shot techniques may be erroneous. This may lead to discarding of network edges that eventually would lead to better network performance.

To address these, we present a novel network pruning method: *Extremely Sparse Pruned Networks* (ESPN). Our method resembles SNIP, but learns the sparse masking operator,  $\mathbf{c}$ , via a standard iterative gradient update framework, instead of using a single-shot estimator. This modification to the standard SNIP framework is also inspired by the observations from Zhou *et al.* [ZLLY19] that learning such sparse indicators can be viewed as a natural process concurrent to training a neural network on data.

**Approach.** Our approach consists of three steps:

1. pretraining  $\mathbf{w}$  while freezing  $\mathbf{c} = \mathbf{1}$ ,
2. leveraging a relaxed form of the SNIP saliency objective to train  $\mathbf{c}$ , thereby pruning the network to required sparsity, and finally,
3. finetuning the pruned network to boost final performance.

For the first step, we train our base architecture on the given dataset to ensure a good initialization for the sparse problem. However, we note that unlike previous works [HMD15, GYC16, LDS90, HS93, LWK18, MAV17], which rely on a fully trained network as input, we do not require our network to be trained to convergence. Instead, we

require the network to only be trained for a few iterations. We also point out that in case of preexisting networks, this step can be safely skipped.

We further modify the SNIP objective in order to iteratively learn masks. First, instead of freezing weights and indicators, we simultaneously train them both, thus keeping track of changes in sensitivity values during training. To achieve this, we modify any given architecture to have an additional matrix associated with each weight matrix such that  $\mathbf{C} = \{\mathbf{C}^i = \mathbf{1}^{m \times n} | \mathbf{W}^i \in \mathbb{R}^{m \times n}\}$ . This is similar to the implementation of the auxiliary variable in SNIP. Secondly, we relax the sparsity constraint in Eq. 3.2 to an  $\ell_1$  penalty, so as to be able to make the objective differentiable. The training objective is given by:

$$\min_{\mathbf{w}, \mathbf{c}} L(f(\mathbf{c} \odot \mathbf{w}, \mathbf{x}), \mathbf{y}) + \alpha \|\mathbf{c}\|_1. \quad (4.1)$$

We rely on standard backpropagation (e.g. SGD) to update both  $\mathbf{w}$  and  $\mathbf{c}$ , until we achieve the required sparsity. We propose a simple update for  $\mathbf{c}$  with  $\mathbf{c} \in \mathbb{R}^m$  initialized to 1 rather than randomly as in [BFG18, LWK18, XWR19, SSM19, SWR20, ZLLY19]. This is advised by Zhou *et al.*'s [ZLLY19] observations regarding the masking operation. Note that optimizing Eq. 4.1 with respect to  $\mathbf{c}$ ,  $\mathbf{c}$  may no longer be sparse with  $c_i \notin \{0, 1\}$ . Assuming that  $\mathbf{c}$  is the optimal selection with salient connections, we update  $\mathbf{w}$  via the element-wise product of  $\mathbf{w}$  and  $\mathbf{c}$  as per Eq. 4.1. Subsequently, we restore  $\mathbf{c}$  to be an indicator function by thresholding,  $\mathbb{1}(\mathbf{c} > \epsilon)$  where  $\epsilon$  is a hyperparameter corresponding to non-zero elements in  $\mathbf{c}$ .

For our algorithm, the choice of the termination condition is significant. Given a target pruning ratio  $p$ , we train both weights  $\mathbf{w}$  and  $\mathbf{c}$  until sparsity of  $\mathbf{c}$  is less or equal to target sparsity  $1 - p$ . While the alternative approach is to train  $\mathbf{w}$  and  $\mathbf{c}$  with a given fixed training budget, the algorithm may either not reach the required sparsity level or may unnecessarily waste computation. Our terminating condition allows us to not only terminate the training, but also removes a sensitive hyperparameter (no. of epochs).

For the third (finetuning) step, we consider two variants. The first variant simply trains the pruned network with the given dataset with a low learning rate until we achieve the desired accuracy (we call this ESPN-FINETUNE; see Alg. 1).

Alternatively, we can also use the ‘rewinding’ technique from [FC18] and [RFC20]. Rewinding involves training the pruned architecture by initializing weights from a previously well-performing supernet. In this case, we use the subset of  $\mathbf{w}_t$  from the warm-up training (trained  $t$  epochs) instead of pretrained weights. After learning auxiliary parameter  $\mathbf{c}$  with a procedure from Alg. 1 Line (4-9), we rewind to epoch  $t$  updating weights by  $\mathbf{w} = \mathbf{w}_t \odot \mathbf{c}$  and train the model with remaining budget. We call this ESPN-REWIND; see Alg. 2.

In the following section, we present a comprehensive evaluation of our approach on various image classification networks and datasets. Our code can be accessed at [https://github.com/chomd90/extreme\\_sparse](https://github.com/chomd90/extreme_sparse).

## 5 Experiments and Results

We analyse the performance of our approach, ESPN, on various image classification tasks. For purposes of our analysis, we consider three architectures, (1) simple fully-connected networks, (2) VGG architectures, and (3) residual networks, for various pruning ratios. Specifically, we first consider fully connected networks trained for MNIST in order to demonstrate the efficacy of our approach. Subsequently, we study the efficacy of ESPN for compressing two massively overparameterized architectures, VGG and ResNet trained on complex datasets such as CIFAR-10/100 [Kri09] and ImageNet [RDS<sup>+</sup>15].

We compare our approach with several pruning methods: (1) SNIP [LAT19], GraSP [WZG20], (2) stabilized Lottery Ticket Hypothesis (LT) [FDRC19] and (3) Dynamic Sparse Reparameterization (DSR) [MW19]. Additionally, we conduct ablation studies to understand the roles of various algorithm components.

**Experimental Setup.** To ensure fair comparisons, we run all algorithms with official implementations and with the best hyperparameters reported in the literature. In the case of ESPN-REWIND, for all experiments except ImageNet, we train the network for 160 epochs through SGD with learning rate 0.1, momentum parameter 0.9, and weight decay 0.0005. We also decay the learning rate with a factor of 0.1 at epochs 80 and 120. For ImageNet, we adapt the official PyTorch implementation [PGM<sup>+</sup>19] to train the pruned network for ESPN-REWIND without modifying the hyperparameter setups (90 epochs, learning rate 0.1, learning rate decay 0.1 every 30 epochs, and weight decay 0.0001).

For ESPN-FINETUNE, we subsequently train the network for 50 epochs with SGD with a learning rate of 0.001 with a decay factor of 0.1 at epoch 30. We also use the weight decay coefficient with 0.0005 for all training except ImageNet. We use the pretrained ResNet50 from Pytorch library. For ESPN finetuning stage, we train ResNet50 on ImageNet with 2/3 of training epochs to the official pytorch implementation (60 epoch, learning rate 0.01, learning rate decay

Table 1: **LeNet300 and LeNet5-Caffe Comparison on MNIST and Fashion-MNIST.** Note that ESPN outperforms all other approaches for various pruning ratios except on LeNet5-Caffe trained with Fashion-MNIST dataset pruning 95% ratio.

Dataset	MNIST				Fashion-MNIST			
<b>LeNet300-100</b>	Acc: 98.80%		Params: 266K		Acc: 89.81%		Params: 266K	
<b>Pruning Ratio</b>	95%	98%	99%	99.6%	95%	98%	99%	99.6%
	(13K)	(5.3K)	(2.7K)	(1.1K)	(13K)	(5.3K)	(2.7K)	(1.1K)
SNIP	97.86	97.15	95.27	86.57	88.31	87.14	81.93	68.60
GraSP	97.96	96.96	93.54	45.02	88.47	86.59	77.62	38.37
LT <sup>+</sup> (Rewind)	98.26	97.74	96.95	92.90	89.55	88.59	87.38	83.57
ESPN-Rewind	<b>98.57</b>	<b>98.39</b>	97.94	97.24	<b>89.94</b>	<b>89.33</b>	<b>88.87</b>	<b>87.74</b>
ESPN-Finetune	98.52	98.35	<b>98.24</b>	<b>97.28</b>	89.59	88.53	88.16	87.67
<b>LeNet5-Caffe</b>	Acc: 99.32%		Params: 431K		Acc: 90.48%		Params: 431K	
<b>Pruning Ratio</b>	95%	98%	99%	99.6%	95%	98%	99%	99.6%
	(22K)	(8.6K)	(4.3K)	(1.7K)	(22K)	(8.6K)	(4.3K)	(1.7K)
SNIP	99.33	99.06	98.93	97.24	90.89	90.30	89.69	84.35
GraSP	99.26	99.21	98.47	97.08	90.58	90.43	89.25	85.79
LT <sup>+</sup> (Rewind)	99.27	<b>99.27</b>	99.17	98.30	<b>91.57</b>	91.11	90.36	87.96
ESPN-Rewind	<b>99.37</b>	99.26	<b>99.22</b>	99.06	91.48	91.15	91.20	90.37
ESPN-Finetune	99.30	99.26	99.10	<b>99.10</b>	91.22	<b>91.48</b>	<b>91.60</b>	<b>90.94</b>

<sup>+</sup> LT rewinded to epoch 1 after magnitude-based pruning on fully-trained networks.

0.1 at 30 and 50 epoch). For LT, we prune the fully-trained network with respect to magnitude and rewind to the early epoch as suggested in [RFC20].

While we train both model weights and auxiliary parameters, we use a standard SGD with Nesterov-momentum 0.9, and no weight decay penalty.

## 5.1 Experimental Results




Figure 1: Sparsity-accuracy tradeoff curve on ResNet32 with CIFAR10/100 dataset. Tested on extreme pruning ratios: {90%, 95%, 98%, 99%, 99.5%, 99.8%}. 99.8% pruned ResNet32 has only 3.8k non-zero parameters. Note that ESPN-finetuning outperforms all other methods for extremely high sparsity levels while being comparable at lower sparsity levels.

**MNIST and Fashion-MNIST Dataset.** We start by evaluating ESPN for LeNet300 and LeNet5-Caffe. For the purposes of comparison, we use SNIP and GraSP as baselines. LeNet300 and LeNet5-Caffe consist of 266K and 431K parameters respectively. We observe that the performance of ESPN to 95% pruning ratio is at par or better when compared to other approaches as shown in Table 1. Upon more severe pruning (99% and 99.6%), ESPN outperforms

Table 2: **VGG19 and ResNet32: Comparison on CIFAR10 and CIFAR100.** ESPN compares favorably with DSR for the case of VGG-19 and CIFAR-10 while being better for the harder problem of CIFAR-100. For the more complicated ResNet32 model, ESPN outperforms other methods across all settings. DSR also fails to converge for extremely sparse settings ( $> 99\%$ ) in the case of ResNet model whereas our method significantly outperforms every other approach.

Dataset	CIFAR10				CIFAR100			
<b>VGG19</b>	Acc: 93.53%		Params: 20M		Acc: 73.96%		Params: 20M	
<b>Pruning Ratio</b>	95%	98%	99%	99.5%	95%	98%	99%	99.5%
	(1M)	(400K)	(200K)	(100K)	(1M)	(400K)	(200K)	(100K)
SNIP [LAT19]	92.97	92.37	10.00 <sup>#</sup>	10.00 <sup>#</sup>	71.90	19.60	1.00 <sup>#</sup>	1.00 <sup>#</sup>
GraSP[WZG20]	92.81	91.94	91.27	88.62	71.28	68.72	65.84	60.28
DSR[MW19]	<b>94.00</b>	<b>93.57</b>	<b>93.15</b>	91.62	<b>72.96</b>	70.77	69.70	66.79
LT <sup>+</sup> [FDRC19]	93.15	92.70	91.29	10.00 <sup>#</sup>	70.97	69.13	66.32	17.60
ESPN-Rewind	93.57	92.72	91.88	<b>91.94</b>	71.68	<b>70.85</b>	69.48	<b>67.93</b>
ESPN-Finetune	93.62	93.24	92.87	91.88	72.32	71.00	<b>70.35</b>	66.45
<b>ResNet32</b>	Acc: 93.93%		Params: 1.9M		Acc: 74.83%		Params: 1.9M	
<b>Pruning Ratio</b>	95%	98%	99%	99.5%	95%	98%	99%	99.5%
	(95K)	(38K)	(19K)	(9.5K)	(95K)	(38K)	(19K)	(9.5K)
SNIP [LAT19]	91.20	88.31	83.35	78.36	63.82	54.09	38.32	27.38
GraSP [WZG20]	91.69	89.01	85.06	80.46	66.20	56.90	47.30	32.63
DSR [MW19]	92.80	90.46	44.96	41.86	69.46	63.56	12.84 <sup>*</sup>	8.11 <sup>*</sup>
LT <sup>+</sup> [FDRC19]	90.57	88.51	85.81	80.31	65.92	57.62	48.74	36.22
ESPN-Rewind	91.83	90.54	89.93	<b>89.31</b>	70.76	69.42	64.83	56.88
ESPN-Finetune	<b>93.06</b>	<b>92.49</b>	<b>90.65</b>	88.77	<b>73.28</b>	<b>70.35</b>	<b>64.89</b>	<b>59.91</b>

<sup>#</sup> Failed to converge with accuracy 1/(number of classes).

<sup>\*</sup> DSR failed to converge.

<sup>+</sup> LT rewinded to a epoch 5 after magnitude-based pruning on fully-trained networks.

all SNIP, GraSP, and LT with only minor degradation in test accuracy. Surprisingly, ESPN-Finetune finds the “lottery ticket” on LeNet5-Caffe with Fashion-MNIST at **99.6%** pruning ratio while achieving **90.94%** test accuracy.

**CIFAR10/100 and Tiny-ImageNet/ImageNet Dataset.** We now evaluate ESPN on modern architectures, VGG19 and ResNet32 on CIFAR10/100, and Tiny-ImageNet image classification datasets. Note that the total number of parameters of VGG19 and ResNet32 are 20M and 1.9M respectively.<sup>2</sup> For VGG19 with CIFAR10, we show comparable performance with DSR (the current state-of-the-art) for lower compression ratios. However, we outperform all other existing algorithms. Specifically, we draw attention to the high pruning ratio of 99.5%, where we report minimal degradation of accuracy for a highly sparse network with only  $\sim 100k$  parameters. We observe that ESPN outperforms SNIP, GraSP, LT, and DSR for all other cases except VGG19 with CIFAR10 shown in Table 2. Especially, all our candidate algorithms except ESPN faces huge degradation in performance when pruning extreme pruning ratio (99% and 99.5%). We notice that ESPN-Finetune represents the Pareto frontiers for ResNet32 on high pruning ratios: {90%, 95%, 98%, 99%, 99.5%, 99.8%} as shown in Figure 1. Overall, we observe that ESPN learns meaningful subnetworks with empirical evidence of achieving extreme compression ratios with minor accuracy degradation.

We observe similar performance on Tiny-ImageNet as that on CIFAR10/100 as seen in Table 3. While DSR achieved significantly higher test accuracy on 90%-pruned VGG19 with CIFAR10, ESPN performs on par with DSR at higher pruning levels. Similar to the CIFAR10/100 case, ESPN outperforms all other candidate algorithms for ResNet32 for three different pruning ratios. We note that for the highest pruning ratio (98%), ESPN outperforms other approaches by a large margin.

We further test our algorithm on ResNet50 (25.6M parameters) for the ImageNet dataset. We test two different pruning ratios {80%, 90%} using our approach and compare with reported results for SNIP, GraSP, and DSR. Note that our approach surpasses SNIP and GraSP for all pruning ratios, while being comparable to DSR. Specifically, ESPN-FINETUNE outperforms DSR in *top-1* accuracy for the 80% case, while being comparable in all other cases.

<sup>2</sup>We use the ResNet architecture defined in [WZG20] for our analysis.

Table 3: **VGG19 and ResNet32: Comparison on Tiny-Imagenet.** ESPN favorably surpasses state of the art approaches on most settings. Note that ESPN and DSR, being dynamic reparameterization methods are proven to exceed single-shot performance. ESPN also exceeds the performance of LT by a significant margin.

Architecture	VGG19: 61.70% (20M)			ResNet32: 61.60% (1.9M)		
Pruning Ratio	90% (2M)	95% (1M)	98% (400K)	90% (190K)	95% (95K)	98% (38K)
SNIP [LAT19]	61.07	57.12	0.5	51.56	40.41	24.81
GraSP [WZG20]	60.26	59.53	56.54	54.84	48.45	37.25
DSR [MW19]	<b>62.43*</b>	59.81*	58.36*	57.19*	56.08*	12.42*
LT+ [FDRC19]	60.69	59.28	56.59	55.64	50.13	39.90
ESPN-Rewind	60.05	<b>59.86</b>	<b>58.66</b>	58.48	57.60	54.21
ESPN-Finetune	59.46	59.27	57.67	<b>60.08</b>	<b>59.83</b>	<b>54.56</b>

\* Experimental results from [WZG20].

\* DSR failed to converge.

+ LT rewinded to Epoch 5 after magnitude-based pruning on fully-trained networks.

Table 4: **ResNet50 on Imagenet**

Architecture	ResNet50 (25.6m)			
Pruning Ratio ( $\kappa$ )	80% (7.3m)		90% (5.1m)	
Test Accuracy	Top-1	Top-5	Top-1	Top-5
Unpruned Model	76.15	92.87	-	-
SNIP*	69.67	89.24	61.97	82.90
GraSP*	72.06	90.82	68.14	88.67
DSR#	73.3	92.4	71.6	90.5
ESPN-Rewind	72.60	91.08	68.70	89.00
ESPN-Finetune	<b>74.34</b>	92.10	71.35	<b>90.68</b>

\* Experimental results from [WZG20].

# Experimental results from [MW19].

## 5.2 Visualizing Weight Distribution of Pruned Networks

To understand better why ESPN outperforms other existing pruning approach in extreme pruning ratio, we visualize sparsity ratios for each layer of the network. We analyse VGG19 and ResNet32 trained with CIFAR10 for two pruning ratios:  $p = 90\%$  and  $98\%$ . Conventional pruning algorithms tend to remove fewer weights in earlier layers (to preserve fine features of the input) and prune more in the deeper layers with higher number of parameters. While VGG19 weight distributions from ESPN follows this trend, we observe that ESPN shows a different trend on ResNet32 compared to SNIP and Lottery Ticket hypothesis (Fig. 2). SNIP follows trend of conventional methods by pruning deeper layers aggressively while preserving weights in the beginning. LT’s weight distribution is comparatively uniform across layers. Given that LT performs better than SNIP (refer Sec. 5.1), we hypothesize that pruning deeper layers in ResNet32 aggressively may induce information bottlenecks, degrading the performance significantly. ESPN, counter-intuitively, prunes more in the earlier layers and the middle layers than in deeper layers. Considering that ESPN outperforms the SNIP and LT for most of the case, ESPN weight distribution counters the conventional pruning intuition and emphasizes the importance of careful rather than excessive pruning in deeper layers.

In case of VGG19, we observe opposite trend to ResNet32 on ESPN weight distribution. While ESPN-Finetune prunes more in the earlier layer and the middle layers than in deeper layers in ResNet32, ESPN aggressively prunes deeper layer compare to SNIP and LT. The presented weight distributions show that the pruning strategy using the same algorithm can differ based on the network architectures.

## 5.3 Ablation Studies

**Role of weight updates, auxiliary parameters, and the L1 penalty.** In the previous section, we have shown that our approach can prune the neural networks with various pruning ratios, specifically for extreme cases ( $> 99\%$ ) with minor sparsity-accuracy tradeoff. We now analyse each of the components of ESPN through ablation studies.



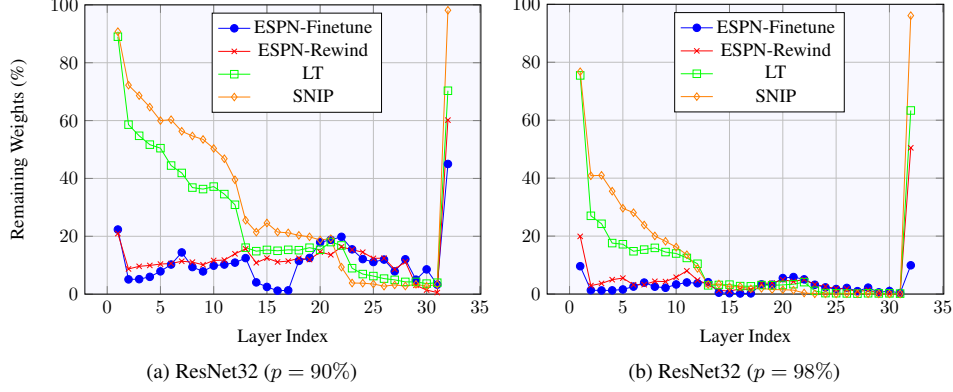


Figure 2: Weight distribution comparisons on ResNet32 with CIFAR10. Almost all pruning algorithms (including ESPN) tend to prune weights in the middle layers. Surprisingly though, ESPN also tends to prune initial layers, concentrating most non-zero weights in the final layers. We hypothesize that learning masks allows us to specifically learn sparser abstractions in earlier layers, analogous to learning sparse features in older classification techniques.

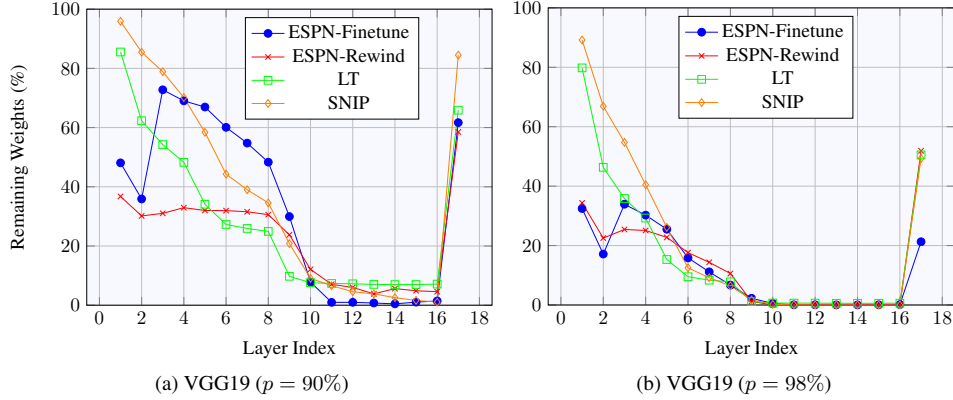


Figure 3: Weight distribution comparisons on VGG19 with CIFAR10. ESPN prunes aggressively in the first two layers and deeper layers. Considering ESPN outperforms SNIP, GraSP, and LT for the most of the case, good pruning strategies may be different depending on the architectures as ESPN’s VGG19 weight distribution trend is opposite to the results from ResNet32.

We consider four different scenarios while learning the auxiliary parameter  $c$ : (1) updating both  $w$  and  $c$  with L1 penalty on  $c$  (original ESPN), (2) only updating  $c$  with L1 penalty on  $c$ , (3) updating  $w$  and  $c$  without L1 penalty, and (4) only updating  $c$  without L1 penalty. Then we compare the test accuracy after the fine-tuning the model (Line 13). We experiment on VGG19 with CIFAR10/100 datasets with pruning ratio  $\{70\%, 80\%, 90\%, 95\%, 95\%, 99\%, 99.5\%, 99.8\%, 99.9\%\}$  which includes common to extreme pruning ratio. We observe that freezing weights and optimizing for only  $c$  shows similar performance as ESPN, but is unstable for some pruning ratios. We also see improvement in performance when both  $w$  and  $c$  are updated. The results are shown in Fig. 4

#### 5.4 Stability on Lasso Coefficient and Learning rate

We check our algorithm’s stability on lasso coefficient  $\alpha$  and learning rate  $\eta$  by observing the sparsity of  $c$  while learning the auxiliary parameters. We conduct experiments on tracking non-zero elements in  $c$  with various learning rate and lasso coefficient. We use pretrained LeNet300 (266K parameters) with MNIST dataset. Our experiment shows that both strength and rate of shrinkage on  $c$  are proportional to learning rate ( $\eta$ ) and lasso coefficient  $\alpha$ . This shows that an extreme pruning ratio requires higher lasso coefficients to meet the condition sparsity of  $c$  less or equal to a targeted number of weights.

An interesting observation is that overall sparsity ratios post training are nearly independent of the learning rate, showing that the network is generally robust to reasonable hyperparameter choices. However, in the case of  $\ell_1$  penalty, the choice of the coefficient needs to be high enough to ensure that we can achieve the required high sparsity.

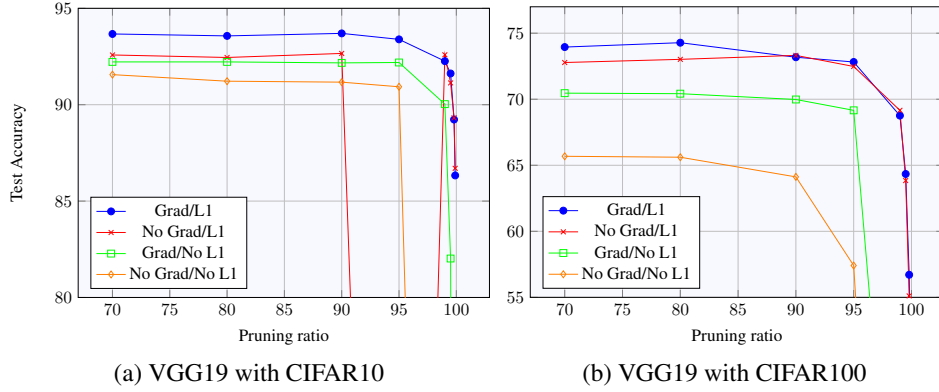


Figure 4: Ablation study on four different setups learning the auxiliary parameters: “Grad/No Grad” and “L1/No L1” correspond to weights updated or not and L1 penalty on auxiliary parameters or not, respectively in stage 1. We test from regular to extreme pruning ratio: {70%, 80%, 90%, 95%, 99%, 99.5%, 99.8%, 99.9%}.

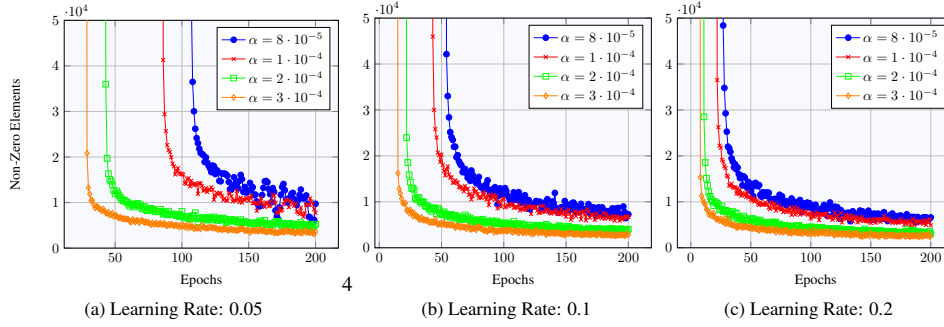


Figure 5: Strength and rate of shrinkage comparison depending on lasso coefficient and learning rate with LeNet300.

## 6 Discussion and Conclusions

In this work, we provide a new algorithm, ESPN, which is a simple and scalable approach to prune a variety of neural network models. While ESPN achieves comparable (even improved) accuracy to SNIP, our algorithm is successfully able to compress the network to extremely high pruning levels ( $> 99\%$ ), up to the regime where the number of parameters are comparable to the input size. To the best of our knowledge, our approach is the first to achieve such high compression ratios for large networks such as ResNet32.

There still exist several open directions for further research. First, ESPN can perhaps be further expanded to a structured-pruning setup that enforces structural sparsity at level of either neurons, or convolutional filters, using group-lasso type constraints [YL07]; this may provide the extra advantage of faster inference over unstructured pruning. Secondly, analysing ESPN’s performance on large-scale language models such as GPT [RWC<sup>+</sup>19]/BERT [DCLT19] and other transformer models would be an important study to undertake. Finally, it would be interesting to carefully understand why our compressed models generalize so well, given that their high levels of compression.

## References

- [AHS17] Sajid Anwar, Kyuyeon Hwang, and Wonyong Sung, *Structured pruning of deep convolutional neural networks*, ACM Journal on Emerging Technologies in Computing Systems (JETC) **13** (2017), no. 3, 1–18.
- [BFG18] Vasileios Belagiannis, Azade Farshad, and Fabio Galasso, *Adversarial network compression*, Euro. Conf. Comp. Vision, 2018, pp. 0–0.
- [BKML18] Guillaume Bellec, David Kappel, Wolfgang Maass, and Robert Legenstein, *Deep rewiring: Training very sparse deep networks*, Proc. Int. Conf. Learning Representations (ICLR), 2018.

- [Cha89] Yves Chauvin, *A back-propagation algorithm with optimal use of hidden units*, Adv. Neural Inf. Proc. Sys. (NeurIPS) (D. S. Touretzky, ed.), Morgan-Kaufmann, 1989, pp. 519–526.
- [CK14] Maxwell D Collins and Pushmeet Kohli, *Memory bounded deep convolutional networks*, arXiv preprint arXiv:1412.1442 (2014).
- [CPI18] Miguel A Carreira-Perpinán and Yerlan Idelbayev, “*learning-compression*” *algorithms for neural net pruning*, IEEE Conf. Comp. Vision and Pattern Recog, 2018, pp. 8532–8541.
- [CWT<sup>+</sup>15] Wenlin Chen, James T. Wilson, Stephen Tyree, Kilian Q. Weinberger, and Yixin Chen, *Compressing neural networks with the hashing trick*, Proc. Int. Conf. Machine Learning (ICML), 2015.
- [CWZZ17] Yu Cheng, Duo Wang, Pan Zhou, and Tao Zhang, *A survey of model compression and acceleration for deep neural networks*, arXiv preprint arXiv:1710.09282 (2017).
- [DCLT19] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova, *Bert: Pre-training of deep bidirectional transformers for language understanding*, arXiv preprint, arXiv:1810.04805 (2019).
- [DSD<sup>+</sup>13] Misha Denil, Babak Shakibi, Laurent Dinh, Marc’Aurelio Ranzato, and Nando De Freitas, *Predicting parameters in deep learning*, Adv. Neural Inf. Proc. Sys. (NeurIPS), 2013.
- [DZW18] Bin Dai, Chen Zhu, and David Wipf, *Compressing neural networks using the variational information bottleneck*, Proc. Int. Conf. Machine Learning (ICML), 2018.
- [FC18] Jonathan Frankle and Michael Carbin, *The lottery ticket hypothesis: Finding sparse, trainable neural networks*, Proc. Int. Conf. Learning Representations (ICLR), 2018.
- [FDRC19] Jonathan Frankle, Gintare Karolina Dziugaite, Daniel M. Roy, and Michael Carbin, *Stabilizing the lottery ticket hypothesis*, arXiv preprint arXiv:1903.01611 (2019).
- [GRK17] Scott Gray, Alec Radford, and Diederik P Kingma, *Gpu kernels for block-sparse weights*, arXiv preprint arXiv:1711.09224 (2017).
- [GYC16] Yiwen Guo, Anbang Yao, and Yurong Chen, *Dynamic network surgery for efficient dnns*, 2016.
- [HCS<sup>+</sup>17] Itay Hubara, Matthieu Courbariaux, Daniel Soudry, Ran El-Yaniv, and Yoshua Bengio, *Quantized neural networks: Training neural networks with low precision weights and activations*, J. Machine Learning Research **18** (2017), no. 1, 6869–6898.
- [HLL<sup>+</sup>18] Yihui He, Ji Lin, Zhijian Liu, Hanrui Wang, Li-Jia Li, and Song Han, *Amc: Automl for model compression and acceleration on mobile devices*, Euro. Conf. Comp. Vision, 2018, pp. 784–800.
- [HMD15] Song Han, Huizi Mao, and William J Dally, *Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding*, Proc. Int. Conf. Learning Representations (ICLR), 2015.
- [HS93] Babak Hassibi and David G Stork, *Second order derivatives for network pruning: Optimal brain surgeon*, Adv. Neural Inf. Proc. Sys. (NeurIPS), 1993.
- [Kri09] Alex Krizhevsky, *Learning multiple layers of features from tiny images*, Tech. report, 2009.
- [LAT19] Namhoon Lee, Thalaiyasingam Ajanthan, and Philip H. S. Torr, *Snip: Single-shot network pruning based on connection sensitivity*, Proc. Int. Conf. Learning Representations (ICLR), vol. abs/1810.02340, 2019.
- [LDS90] Yann LeCun, John S Denker, and Sara A Solla, *Optimal brain damage*, Adv. Neural Inf. Proc. Sys. (NeurIPS), 1990.
- [LLS<sup>+</sup>17] Zhuang Liu, Jianguo Li, Zhiqiang Shen, Gao Huang, Shoumeng Yan, and Changshui Zhang, *Learning efficient convolutional networks through network slimming*, Proc. IEEE Intl. Conf. Comp. Vision, 2017, pp. 2736–2744.
- [LRLZ17] Ji Lin, Yongming Rao, Jiwen Lu, and Jie Zhou, *Runtime neural pruning*, Adv. Neural Inf. Proc. Sys. (NeurIPS), 2017, pp. 2181–2191.
- [LSZ<sup>+</sup>19] Zhuang Liu, Mingjie Sun, Tinghui Zhou, Gao Huang, and Trevor Darrell, *Rethinking the value of network pruning*, 2019.
- [LUW17] Christos Louizos, Karen Ullrich, and Max Welling, *Bayesian compression for deep learning*, Adv. Neural Inf. Proc. Sys. (NeurIPS), 2017.
- [LWK18] Christos Louizos, Max Welling, and Diederik P. Kingma, *Learning sparse neural networks through l0 regularization*, 2018.
- [LWL17] Jian-Hao Luo, Jianxin Wu, and Weiyao Lin, *Thinet: A filter level pruning method for deep neural network compression*, Proc. IEEE Intl. Conf. Comp. Vision, 2017, pp. 5058–5066.

- [MAV17] Dmitry Molchanov, Arsenii Ashukha, and Dmitry Vetrov, *Variational dropout sparsifies deep neural networks*, Proc. Int. Conf. Machine Learning (ICML), 2017.
- [MMS<sup>+</sup>18] 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** (2018), no. 1, 1–12.
- [MS89] Michael C Mozer and Paul Smolensky, *Skeletonization: A technique for trimming the fat from a network via relevance assessment*, Adv. Neural Inf. Proc. Sys. (NeurIPS) (D. S. Touretzky, ed.), Morgan-Kaufmann, 1989, pp. 107–115.
- [MTK<sup>+</sup>17] Pavlo Molchanov, Stephen Tyree, Tero Karras, Timo Aila, and Jan Kautz, *Pruning convolutional neural networks for resource efficient transfer learning*, Proc. Int. Conf. Learning Representations (ICLR), 2017.
- [MW19] Hesham Mostafa and Xin Wang, *Parameter efficient training of deep convolutional neural networks by dynamic sparse reparameterization*, Proc. Int. Conf. Machine Learning (ICML), 2019.
- [OBSS20] Deniz Oktay, Johannes Ballé, Saurabh Singh, and Abhinav Shrivastava, *Model compression by entropy penalized reparameterization*, 2020.
- [PGM<sup>+</sup>19] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala, *Pytorch: An imperative style, high-performance deep learning library*, Adv. Neural Inf. Proc. Sys. (NeurIPS), 2019.
- [PPA18] Antonio Polino, Razvan Pascanu, and Dan Alistarh, *Model compression via distillation and quantization*, Proc. Int. Conf. Learning Representations (ICLR), 2018.
- [RDS<sup>+</sup>15] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei, *ImageNet Large Scale Visual Recognition Challenge*, Intl. J. Comp. Vision **115** (2015), no. 3, 211–252.
- [RFC20] Alex Renda, Jonathan Frankle, and Michael Carbin, *Comparing rewinding and fine-tuning in neural network pruning*, Proc. Int. Conf. Learning Representations (ICLR) (2020).
- [RWC<sup>+</sup>19] Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever, *Language models are unsupervised multitask learners*.
- [SB16] Bharat Bhusan Sau and Vineeth N Balasubramanian, *Deep model compression: Distilling knowledge from noisy teachers*, arXiv preprint arXiv:1610.09650 (2016).
- [Set97] Rudy Setiono, *A penalty-function approach for pruning feedforward neural networks*, Neural computation **9** (1997), no. 1, 185–204.
- [SSM19] Pedro Savarese, Hugo Silva, and Michael Maire, *Winning the lottery with continuous sparsification*, arXiv preprint arXiv:1912.04427 (2019).
- [SWR20] Victor Sanh, Thomas Wolf, and Alexander M Rush, *Movement pruning: Adaptive sparsity by fine-tuning*, arXiv preprint arXiv:2005.07683 (2020).
- [TKTH18] Lucas Theis, Iryna Korshunova, Alykhan Tejani, and Ferenc Huszár, *Faster gaze prediction with dense networks and fisher pruning*, arXiv preprint arXiv:1801.05787 (2018).
- [TLFF18] Enzo Tartaglione, Skjalg Lepsøy, Attilio Fiandrotti, and Gianluca Francini, *Learning sparse neural networks via sensitivity-driven regularization*, Adv. Neural Inf. Proc. Sys. (NeurIPS), 2018, pp. 3878–3888.
- [TVDJ20] Hugo Touvron, Andrea Vedaldi, Matthijs Douze, and Hervé Jégou, *Fixing the train-test resolution discrepancy: Fixefficientnet*, arXiv preprint, arXiv:2003.08237 (2020).
- [UMW17] Karen Ullrich, Edward Meeds, and Max Welling, *Soft weight-sharing for neural network compression*, Proc. Int. Conf. Learning Representations (ICLR), 2017.
- [VSF20] Stijn Verdenius, Maarten Stol, and Patrick Forré, *Pruning via iterative ranking of sensitivity statistics*, arXiv preprint arXiv:2006.00896 (2020).
- [WZG20] Chaoqi Wang, Guodong Zhang, and Roger Grosse, *Picking winning tickets before training by preserving gradient flow*, Proc. Int. Conf. Learning Representations (ICLR) (2020).
- [XWR19] Xia Xiao, Zigeng Wang, and Sanguthevar Rajasekaran, *Autoprune: Automatic network pruning by regularizing auxiliary parameters*, Adv. Neural Inf. Proc. Sys. (NeurIPS), 2019.
- [YL07] Ming Yuan and Yi Lin, *Model selection and estimation in the gaussian graphical model*, Biometrika **94** (2007), no. 1, 19–35.

- [YMD<sup>+</sup>15] Zichao Yang, Marcin Moczulski, Misha Denil, Nando de Freitas, Alexander J. Smola, Le Song, and Ziyu Wang, *Deep fried convnets*, Proc. IEEE Intl. Conf. Comp. Vision, 2015.
- [ZG17] Michael Zhu and Suyog Gupta, *To prune, or not to prune: exploring the efficacy of pruning for model compression*, arXiv preprint arXiv:1710.01878 (2017).
- [ZLLY19] Hattie Zhou, Janice Lan, Rosanne Liu, and Jason Yosinski, *Deconstructing lottery tickets: Zeros, signs, and the supermask*, Adv. Neural Inf. Proc. Sys. (NeurIPS), 2019, pp. 3592–3602.