

---

# Parameter Efficient Training of Deep Convolutional Neural Networks by Dynamic Sparse Reparameterization

---

Hesham Mostafa<sup>1</sup> Xin Wang<sup>1,2</sup>

## Abstract

Modern deep neural networks are typically highly overparameterized. Pruning techniques are able to remove a significant fraction of network parameters with little loss in accuracy. Recently, techniques based on dynamic reallocation of non-zero parameters have emerged, allowing direct training of sparse networks without having to pre-train a large dense model. Here we present a novel dynamic sparse reparameterization method that addresses the limitations of previous techniques such as high computational cost and the need for manual configuration of the number of free parameters allocated to each layer. We evaluate the performance of dynamic reallocation methods in training deep convolutional networks and show that our method outperforms previous static and dynamic reparameterization methods, yielding the best accuracy for a fixed parameter budget, on par with accuracies obtained by iteratively pruning a pre-trained dense model. We further investigated the mechanisms underlying the superior generalization performance of the resultant sparse networks. We found that neither the structure, nor the initialization of the non-zero parameters were sufficient to explain the superior performance. Rather, effective learning crucially depended on the continuous exploration of the sparse network structure space during training. Our work suggests that exploring structural degrees of freedom during training is more effective than adding extra parameters to the network.

---

<sup>1</sup>Artificial Intelligence Products Group, Intel Corporation, San Diego, CA, USA. <sup>2</sup>Currently with Cerebras Systems, Los Altos, CA, USA, work done while at Intel Corporation. Correspondence to: Xin Wang <xin@cerebras.net>.

## 1. Introduction

Deep neural networks' success in a wide range of application domains, ranging from computer vision to machine translation to automatic speech recognition, stems from their ability to learn complex transformations by data examples while achieving superior generalization performance. Though they generalize well, deep networks learn more effectively when they are highly overparameterized (Brutzkus et al., 2017; Zhang et al., 2016). Emerging evidence has attributed this need for overparameterization to the geometry of the high-dimensional loss landscapes (Dauphin et al., 2014; Choromanska et al., 2014; Goodfellow et al., 2014; Im et al., 2016; Wu et al., 2017; Liao & Poggio, 2017; Cooper, 2018; Novak et al., 2018), and to the implicit regularization properties of stochastic gradient descent (SGD) (Brutzkus et al., 2017; Zhang et al., 2018a; Poggio et al., 2017), though a thorough theoretical understanding is not yet complete.

In practice, multiple techniques are able to compress large trained models, including distillation (Bucilua et al., 2006; Hinton et al., 2015), weight precision reduction (Hubara et al., 2016; McDonnell, 2018), low-rank decomposition (Jaderberg et al., 2014; Denil et al., 2013), and pruning (Han et al., 2015a; Zhang et al., 2018b). While these methods are highly effective in reducing the size of network parameters with little degradation in accuracy, they either operate on a pre-trained model or require the full overparameterized model to be stored and updated during, or at least at a certain stage of, training. Thus, training remains memory-inefficient despite the compact size of the resultant network produced by compression. The effectiveness of these compression methods, however, indicates the existence of compact network parameter configurations that are able to generalize on par with large networks. This raises a tantalizing hypothesis that overparameterization during training might not be a strict necessity and alternative training or reparameterization methods might exist to discover and train compact networks directly.

The problem of achieving training-time parameter efficiency is being approached in a number of ways. Most straightforward is to search for more parameter efficient network architectures. Innovations in this direction for deep convolutional neural networks (CNNs) include adoption of skip connec-

tions (He et al., 2015), replacement of fully-connected layers with global average pooling layers followed directly by the classifier layer (Lin et al., 2013), and depthwise separable convolutions (Sifre & Mallat, 2014; Howard et al., 2017). These modern CNN architectures drastically improved the accuracies achievable at a given parameter budget.

Instead of inventing new network architectures, an alternative approach is to reparameterize an existing model architecture, which is the approach we take in this work. In general, any *differentiable reparameterization* can be used to augment training of a given model. Let an original network (or a layer therein) be denoted by  $\mathbf{y} = f(\mathbf{x}; \boldsymbol{\theta})$ , parameterized by  $\boldsymbol{\theta} \in \Theta$ . Reparameterize it by  $\boldsymbol{\phi} \in \Phi$  and  $\boldsymbol{\psi} \in \Psi$  through  $\boldsymbol{\theta} = g(\boldsymbol{\phi}; \boldsymbol{\psi})$ , where  $g$  is differentiable w.r.t.  $\boldsymbol{\phi}$  but not necessarily w.r.t.  $\boldsymbol{\psi}$ . Denote the reparameterized network by  $f_\psi$ , considering  $\boldsymbol{\psi}$  as *metaparameters*<sup>\*</sup>:

$$\mathbf{y} = f(\mathbf{x}; g(\boldsymbol{\phi}; \boldsymbol{\psi})) \triangleq f_\psi(\mathbf{x}; \boldsymbol{\phi}). \quad (1)$$

$f_\psi$  is trained by backpropagating errors through  $g$ , as  $\frac{\partial}{\partial \boldsymbol{\phi}} = \frac{\partial g}{\partial \boldsymbol{\phi}} \frac{\partial}{\partial g}$ . If it is so chosen that  $\dim(\Phi) + \dim(\Psi) < \dim(\Theta)$  and  $f_\psi \approx f$  in terms of generalization performance, then  $f_\psi$  is a more parameter efficient function approximator than  $f$ .

*Sparse reparameterization* is a special case where  $g$  is a linear projection;  $\boldsymbol{\phi}$  is the non-zero entries (i.e. “weights”) and  $\boldsymbol{\psi}$  their indices (i.e. “connectivity”) in the original parameter tensor  $\boldsymbol{\theta}$ . Likewise, *parameter sharing* is a similar special case of linear reparameterization where  $\boldsymbol{\phi}$  is the tied parameters and  $\boldsymbol{\psi}$  the indices at which each parameter is placed (with repetition) in the original parameter tensor  $\boldsymbol{\theta}$ . If metaparameters  $\boldsymbol{\psi}$  are fixed during the course of training, the reparameterization is *static*, whereas if  $\boldsymbol{\psi}$  is adjusted adaptively during training, we call it *dynamic* reparameterization.

In this paper, we investigate multiple static and dynamic reparameterizations of deep residual CNNs for efficient training. Inspired by previous techniques, we developed a novel dynamic reparameterization method that yielded the highest parameter efficiency in training sparse deep residual networks, outperforming existing static and dynamic reparameterization methods.

Our method dynamically changes the sparse structure of the network during training. Its superior performance suggests that, given a limited storage and computational budget for training a CNN, it is better to allocate part of the resources to describing and evolving the structure of the network, than to spend it entirely on the parameters of a dense network.

<sup>\*</sup>We use the term *metaparameter* to refer to the parameters  $\boldsymbol{\psi}$  of the reparameterization function  $g$ . They differ from parameters  $\boldsymbol{\phi}$  in that they are not optimized through gradient descent, and they differ from hyperparameters in that they define meaningful features of the model which are required for inference.

Furthermore, we show that the success of dynamic sparse reparameterization is not solely due to the final sparse structure of the resultant networks, nor to a combination of final structure and initial weight values. Rather, training-time structural exploration is necessary for best generalization, even if a high-performance structure and its initial values are known *a priori*. Thus, simultaneous exploration of network structure and parameter optimization through gradient descent are synergistic. Structural exploration improves the trainability of sparse deep CNNs.

## 2. Related work

Training of differentially reparameterized networks has been proposed in numerous studies before.

**Dense reparameterization** Several dense reparameterization techniques sought to reduce the size of fully connected layers. These include low-rank decomposition (Denil et al., 2013), fastfood transform (Yang et al., 2014), ACDC transform (Moczulski et al., 2015), HashedNet (Chen et al., 2015), low displacement rank (Sindhwani et al., 2015) and block-circulant matrix parameterization (Treister et al., 2018).

Note that similar reparameterizations were also used to introduce certain algebraic properties to the parameters for purposes other than reducing model sizes, e.g. to make training more stable as in unitary evolution RNNs (Arjovsky et al., 2015) and in weight normalization (Salimans & Kingma, 2016), to inject inductive biases (Thomas et al., 2018), and to alter (Dinh et al., 2017) or to measure (Li et al., 2018) properties of the loss landscape. These dense reparameterization methods are static.

**Sparse reparameterization** Successful training of sparse reparameterized networks usually employs iterative pruning and retraining, e.g. (Han et al., 2015b; Narang et al., 2017; Zhu & Gupta, 2017)<sup>†</sup>. Training typically starts with a large pre-trained model and sparsity is gradually increased by pruning and fine-tuning. Training a small, static, and sparse model *de novo* fares worse than compressing a large dense model to the same sparsity (Zhu & Gupta, 2017).

(Frankle & Carbin, 2018) identified small and sparse subnetworks post-training which, when trained in isolation, reached a similar accuracy as the enclosing big network. They further showed that these subnetworks were sensitive to initialization, and hypothesized that the role of overpa-

<sup>†</sup>Note that these, as well as all other sparse techniques we benchmark against in this paper, impose *non-structured* sparsification on parameter tensors, yielding *sparse* models. There also exist a class of *structured* pruning methods that “sparsify” at channel or layer granularity, e.g. (Luo et al., 2017) and (Huang & Wang, 2017), generating essentially small *dense* models. We describe a full landscape of existing methods in Appendix D.

parameterization is to provide a large number of candidate subnetworks, thereby increasing the likelihood that one of these subnetworks will have the necessary structure and initialization needed for effective learning.

Most closely related to our work are dynamic sparse reparameterization techniques that emerged only recently. Like ours, these methods adaptively alter, by certain heuristic rules, the location of non-zero parameters during training. Sparse evolutionary training (SET) (Mocanu et al., 2018) used magnitude-based pruning and random growth at the end of each training epoch. NeST (Dai et al., 2017; 2018) iteratively grew and pruned parameters and neurons during training; parameter growth was guided by parameter gradient and pruning by parameter magnitude. Deep Rewiring (DeepR) (Bellec et al., 2017) combined dynamic sparse parameterization with stochastic parameter updates for training. These methods were primarily demonstrated with sparsification of fully-connected layers and applied to relatively small and shallow networks. They also required manual configuration of sparsity levels for different layers of the model. The method we propose in this paper is more scalable and computationally efficient than these previous approaches, while achieving better performance on deep CNNs.

### 3. Methods

We sparsely reparameterize the majority of layers in deep CNNs. All sparse parameter tensors are randomly initialized at the same sparsity (i.e. fraction of zeros). During training, free parameters are moved within and across weight tensors every few hundred training iterations, following a two-phase procedure (Algorithm 1): magnitude-based pruning followed by random growth. Throughout training, we always maintain the same total number of non-zero parameters in the network.

Our magnitude-based pruning is based on an adaptive global threshold  $H$  where all sparse weights with absolute values smaller than  $H$  are pruned.  $H$  is adjusted via a setpoint negative feedback loop to maintain approximately (with tolerance  $\delta$ ) a fixed number of pruned/grown parameters  $N_p$  during each reallocation step.

Immediately after removing  $K$  parameters during the pruning phase,  $K$  zero-initialized parameters are redistributed among the sparse parameter tensors, following a heuristic rule: layers having larger fractions of non-zero weights receive proportionally more free parameters (see Algorithm 1). Intuitively, one should allocate more parameters to layers such that training loss is more quickly reduced. This means, to the first order, free parameters should be redistributed to layers whose parameters receive larger loss gradients. If a layer has been heavily pruned, this indicates that, for a large portion of its parameters, the training loss gradients are not

large or consistent enough to counteract the pull towards zero exerted by weight-decay regularization. This layer is therefore to receive a smaller share of new free parameters during the growth phase. The reallocated parameters are randomly placed at zero positions in the target weight tensors. To ensure the numbers of pruned and grown free parameters are exactly the same, we impose extra guards against rounding errors, as well as against special cases where more free parameters are allocated to a tensor than there are zero positions. For simplicity of exposition, we omit these minor details in Algorithm 1. See Appendix A for a more detailed description of the algorithm.

Our algorithm differs from *SET* (Mocanu et al., 2018) in two important aspects. First, instead of pruning a fixed fraction of weights at each reallocation step, we use an adaptive threshold for pruning. Second, we automatically reallocate parameters across layers during training and do not impose a fixed, manually configured, sparsity level on each layer. The first difference leads to reduced computational overhead as it obviates the need for sorting operations, and the second to better performing networks (see Section 4) and the ability to train extremely sparse networks as shown in Appendix F.

We evaluated our method on the deep residual CNNs listed in Table 1, and compared its performance against existing static and dynamic reparameterization methods<sup>‡</sup>. We did not experiment with AlexNet (Krizhevsky et al., 2012) and VGG-style networks (Simonyan & Zisserman, 2014) as their parameter efficiency is inferior to modern residual networks. Such a choice makes the improvement in parameter efficiency achieved by our dynamic sparse training method more practically relevant. Dynamic sparse reparameterization was applied to all weight tensors of convolutional layers (with the exception of downsampling convolutions and the first convolutional layer acting on the input image), while all biases and parameters of normalization layers were kept dense.

At a specific global sparsity<sup>§</sup>  $s$ , we compared our method (*dynamic sparse*) against six baselines:

1. *Full dense*: original large and dense model, with  $N$  free parameters;
2. *Thin dense*: original model with less wide layers, such that it had the same size as *dynamic sparse*;
3. *Static sparse*: original model initialized at sparsity level  $s$  with random sparseness pattern, then trained with connectivity (sparseness pattern) fixed;
4. *Compressed sparse*: sparse model obtained by iteratively pruning and re-training a large and dense pre-trained

<sup>‡</sup>Code is available at <https://github.com/IntelAI/dynamic-reparameterization>.

<sup>§</sup>Global sparsity  $s$  is defined as the overall sparsity of all sparse parameter tensors, not the entire model, which has a small fraction of dense parameters.

**Algorithm 1:** Reallocation of non-zero parameters within and across parameter tensors

---

```

1: for each sparse parameter tensor  $\mathbf{W}_i$  do
2:    $(\mathbf{W}_i, k_i) \leftarrow \text{prune\_by\_threshold}(\mathbf{W}_i, H)$  ▷  $k_i$  is the number of pruned weights
3:    $l_i \leftarrow \text{number\_of\_nonzero\_entries}(\mathbf{W}_i)$  ▷ Number of surviving weights after pruning
4: end for
5:  $(K, L) \leftarrow (\sum_i k_i, \sum_i l_i)$  ▷ Total number of pruned and surviving weights
6:  $H \leftarrow \text{adjust\_pruning\_threshold}(H, K, \delta)$  ▷ Adjust pruning threshold
7: for each sparse parameter tensor  $\mathbf{W}_i$  do
8:    $\mathbf{W}_i \leftarrow \text{grow\_back}(\mathbf{W}_i, \frac{l_i}{L}K)$  ▷ Grow  $\frac{l_i}{L}K$  zero-initialized weights at random in  $\mathbf{W}_i$ 
9: end for

```

---

- model to target sparsity  $s$  (Zhu & Gupta, 2017);
- 5. *DeepR*: sparse model trained by using Deep Rewiring (Bellec et al., 2017);
- 6. *SET*: sparse model trained by using Sparse Evolutionary Training (SET) (Mocanu et al., 2018).

Appendix C compares our method against an additional static (dense) parameterization method based on weight sharing: *HashedNet* (Chen et al., 2015).

Because sparse tensors require storage of both the free parameter values and their locations, we compare models that have the same size in descriptive length, instead of the same number of weights. While the number of bits needed to specify the connectivity is implementation dependent, we assume one bit is used for each position in the weight tensors indicating whether this position is zero or not. A sparse tensor is fully defined by this bit-mask, together with the non-zero entries. This scheme was previously used in CNN accelerators that natively operate on sparse structures (Aimar et al., 2018). For a network with  $N$  32-bit weights in its dense form, a sparse version at sparsity  $s$  has a descriptive length of  $(32s + 1)N$  bits, and is thus equivalent in size to a thinner dense network with  $(s + \frac{1}{32})N$  weights. We use this formula to determine the parameter counts of the *thin dense* baseline, which has  $\frac{N}{32}$  more weights than comparable sparse models.

A recent study (Liu et al., 2018) showed that training small networks *de novo* can almost always match the generalization performance obtained by post-training pruning of larger networks, so long as the small networks were trained for long enough. To address concerns that the superior performance of *dynamic sparse* might be matched by training *thin dense* or *static sparse* networks for more epochs, we always train *thin dense* and *static sparse* baselines for double the number of epochs used to train *dynamic sparse* models.

Note that *compressed sparse* is a compression method that first trains a large dense model and then iteratively prunes it down, whereas *dynamic sparse* and all other baselines maintain a constant small model size throughout training. For *compressed sparse*, we train the large dense model for the

same number of epochs used for our *dynamic sparse*, and then iteratively prune and fine-tune it across additional training epochs. *Compressed sparse* thus trains for more epochs than *dynamic sparse*. See Appendix B for hyperparameters used for all experiments.

## 4. Experimental results

**WRN-28-2 on CIFAR10** We ran experiments on a Wide Resnet model WRN-28-2 (Zagoruyko & Komodakis, 2016) trained to classify CIFAR10 images (see Appendix B for details of implementation). We varied global sparsity levels and evaluated test accuracy of different training methods based on dynamic and static reparameterization. As shown in Figure 1a, *static sparse* and *thin dense* significantly underperformed *compressed sparse* model as expected, whereas our *dynamic sparse* performed on par or slightly better than *compressed sparse* on average. *DeepR* significantly underperformed all other method. Though *SET* was generally on par with *compressed sparse* and *dynamic sparse* at low sparsity levels, it underperformed *dynamic sparse* at high sparsity levels. Even though the statically reparameterized models *static sparse* and *thin dense* were trained for twice the number of epochs, they still failed to reach the accuracy of *dynamic sparse* or *SET*. Note that *thin dense* had even more trainable free parameters than all sparse models (see Section 3).

Further, we inspected the layer-wise sparsity patterns that emerged from automatic parameter reallocation across layers (Algorithm 1) during *dynamic sparse* training. We observed consistent patterns at different sparsity levels: (a) larger parameter tensors tended to be sparser than smaller ones, and (b) deeper layers tended to be sparser than shallower ones. Figure 1b shows a breakdown of the final sparsity levels of different residual blocks at different sparsity levels.

**Resnet-50 on Imagenet** Next, we trained the Resnet-50 bottleneck architecture (He et al., 2015) on Imagenet (see Appendix B for details of implementation). We ran experiments at two global sparsity levels, 0.8 and 0.9 (Table 2). Models obtained by our (*dynamic sparse*) method

**Table 1:** Datasets and models used in experiments

Dataset	CIFAR10	Imagenet
Model	WRN-28-2 (Zagoruyko & Komodakis, 2016)	Resnet-50 (He et al., 2015)
Architecture	C16/3×3 [C16/3×3,C16/3×3]×4 [C64/3×3,C64/3×3]×4 [C128/3×3,C128/3×3]×4 GlobalAvgPool, F10	C64/7×7-2, MaxPool/3×3-2 [C64/1×1, C64/3×3, C256/1×1]×3 [C128/1×1, C128/3×3, C512/1×1]×4 [C256/1×1, C256/3×3, C1024/1×1]×6 [C512/1×1, C512/3×3, C2048/1×1]×3 GlobalAvgPool, F1000
# Parameters	1.5M	25.6M

For brevity architecture specifications omit batch normalization and activations. Pre-activation batch normalization was used in all cases. Convolutional (C) layers are specified with output size and kernel size and Max pooling (MaxPool) layers with kernel size. Brackets enclose residual blocks postfixed with repetition numbers; the downsampling convolution in the first block of a scale group is implied.

outperformed all dynamic and static reparameterization baseline methods, slightly outperforming *compressed sparse* models obtained through post-training compression of a large dense model. In Table 2, we also list two additional representative methods of *structured* pruning (see Appendix D), *ThiNet* (Luo et al., 2017) and *Sparse Structure Selection* (Huang & Wang, 2017), which, consistent with recent criticisms (Liu et al., 2018), underperformed static dense baselines. Similar to *dynamic sparse* WRN-28-2, reliable sparsity patterns across parameter tensors in different layers emerged from dynamic parameter reallocation during training, displaying the same empirical trends described above (Figure 2).

**Computational overhead of dynamic parameter reallocation** We assessed the additional computational cost incurred by dynamic parameter reallocation steps (Algorithm 1) during training, and compared ours with existing dynamic sparse reparameterization techniques, *DeepR* and *SET* (Table 3). Because both *SET* and *dynamic sparse* reallocate parameters only intermittently (every few hundred training iterations), the computational overhead was negligible for the experiments presented here<sup>¶</sup>. *DeepR*, however, requires adding noise to gradient updates as well as reallocating parameters every training iteration, leading to a significantly larger overhead.

**Understanding the effects of dynamic parameter reallocation** Why did dynamic parameter reallocation yield sparse models that generalize better than static models trained *de novo*? To address this question, we investigated

<sup>¶</sup>Because of the rather negligible overhead, the reduced operation count thanks to the elimination of sorting operations did not amount to a substantial improvement over *SET* on GPUs. Our method’s another advantage over *SET* lies in its ability to produce better sparse models and to reallocate free parameters automatically (see Appendix F).

whether our method discovered more trainable sparse network structures, following the reasoning of the recently proposed “lottery ticket” hypothesis (Frankle & Carbin, 2018).

First, we did the following experiment with WRN-28-2 trained on CIFAR10: after training with *dynamic sparse* method, we retained the final network sparseness pattern (i.e. positions of non-zero entries in all sparse parameter tensors), and then randomly re-initialized this network and re-trained with its structure fixed (Figure 3a). It failed to reach comparable accuracy, suggesting that the sparse connectivity discovered by our method is not sufficient to explain the high generalization performance.

Next, we asked whether the particular weight initialization of the sparse network in addition to its sparseness structure led to high accuracies (Frankle & Carbin, 2018). We used the final network structure as described above, and re-initialized it with the exact same initial values used in the original training. As shown in Figure 3a, the combination of final structure and original initialization still fell significantly short of the level of accuracy achieved by training with dynamic parameter reallocation, not significantly different from training the same network with random re-initialization.

For control, we also show the performance of *static sparse* models where the sparse network structure and its initialization were both random (Figure 3a), which, not surprisingly, performed the worst. Similar trends were observed for Resnet-50 trained on Imagenet (Figure 3b). All static networks, in all sparseness pattern and re-initialization configurations, were trained for double the number of epochs used for dynamic training.

These results suggest that the dynamic evolution of the network through parameter reallocation is crucial to effective learning, because the superior generalization performance

## Dynamic sparse reparameterization

**Table 2:** Test accuracy% (top-1, top-5) of Resnet-50 trained on Imagenet

		Final overall sparsity (# Parameters)	0.8 (7.3M)		0.9 (5.1M)		0.0 (25.6M)		
Reparameterization	Static	<i>Thin dense</i>	72.4	90.9	70.7	89.9			
			[-2.5]	[-1.5]	[-4.2]	[-2.5]			
		<i>Static sparse</i>	71.6	90.4	67.8	88.4			
			[-3.3]	[-2.0]	[-7.1]	[-4.0]			
	Dynamic	<i>DeepR</i> (Bellec et al., 2017)	71.7	90.6	70.2	90.0	74.9	92.4	
			[-3.2]	[-1.8]	[-4.7]	[-2.4]			[0.0]
<i>SET</i> (Mocanu et al., 2018)		72.6	91.2	70.4	90.1				
		[-2.3]	[-1.2]	[-4.5]	[-2.3]				
<i>Dynamic sparse</i> (Ours)	<b>73.3</b>	<b>92.4</b>	<b>71.6</b>	<b>90.5</b>					
		<b>[-1.6]</b>	<b>[ 0.0]</b>	<b>[-3.3]</b>	<b>[-1.9]</b>				
Compression	<i>Compressed sparse</i> (Zhu & Gupta, 2017)		73.2	91.5	70.3	90.0			
			[-1.7]	[-0.9]	[-4.6]	[-2.4]			
	<i>ThiNet</i> (Luo et al., 2017)		68.4	88.3	(at 8.7M parameter count)				
			[-4.5]	[-2.8]					
<i>SSS</i> (Huang & Wang, 2017)		71.8	90.8	(at 15.6M parameter count)					
		[-4.3]	[-2.1]						

Numbers in square brackets are differences from the *full dense* baseline. Romanized numbers are results of our experiments, and italicized ones taken directly from the original paper. Performance of two structured pruning methods, *ThiNet* and *Sparse Structure Selection* (SSS), are also listed for comparison (below the double line, see Appendix D for a discussion of their relevance); note the difference in parameter counts.

cannot be solely attributed to the network’s structure, nor to its initialization, nor to a combination of the two.

Finally, to investigate whether the convergence of sparse network structures and that of parameter values had similar time courses, we did extra experiments with WRN-28-2, where at various stages during training, we stopped dynamic parameter reallocation, fixing the network structure while continuing the optimization of parameter values. As shown in Figure 4, dynamic parameter reallocation does not need to be active for the entire course of training, but only for some initial epochs. This suggests the network structure converges faster than the network parameters, which might

be exploited in practice to further reduce computational cost.

## 5. Discussion

In this work, we addressed the following problem: given a small, fixed budget of parameters for a deep residual CNN throughout training time, how to train it to yield the best generalization performance. We showed that training with dynamic parameter reallocation can achieve significantly better accuracies than static reparameterization at the same model size. Dynamic sparse reparameterization techniques have received relatively little attention to date, two existing methods (*SET* and *DeepR*) being applied only to relatively small and shallow networks. We proposed a dynamic parameterization method that adaptively reallocates free parameters across the network based on a simple heuristic during training. Our method yielded sparse models that generalize better than those produced by previous dynamic parameterization methods and outperformed all the static reparameterization methods we benchmarked against<sup>||</sup>.

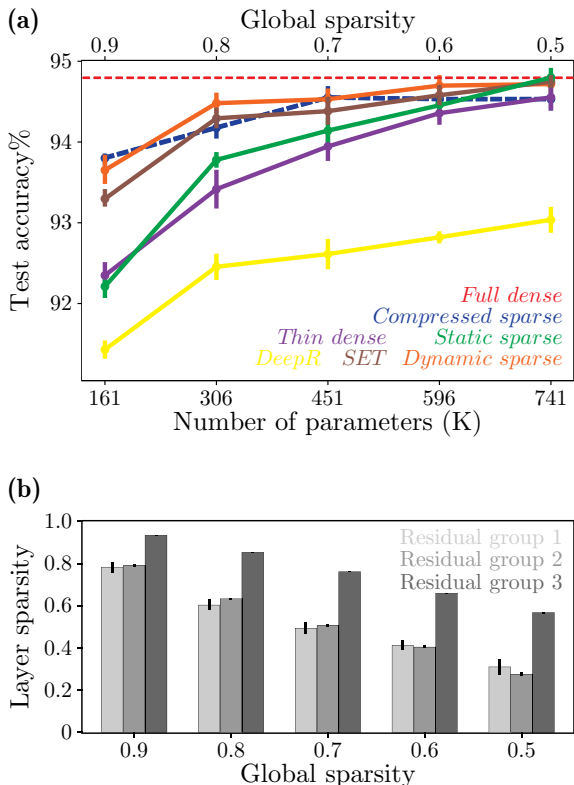
High-performance sparse networks are often obtained by

<sup>||</sup> Additionally, we show that our method outperformed a static dense reparameterization method *HashedNet* (Chen et al., 2015) (Appendix C), and that it is also able to train networks at extreme sparsity levels where previous static and dynamic parameterization methods often fail catastrophically (Appendix F).

**Table 3:** Wall-clock training time comparison

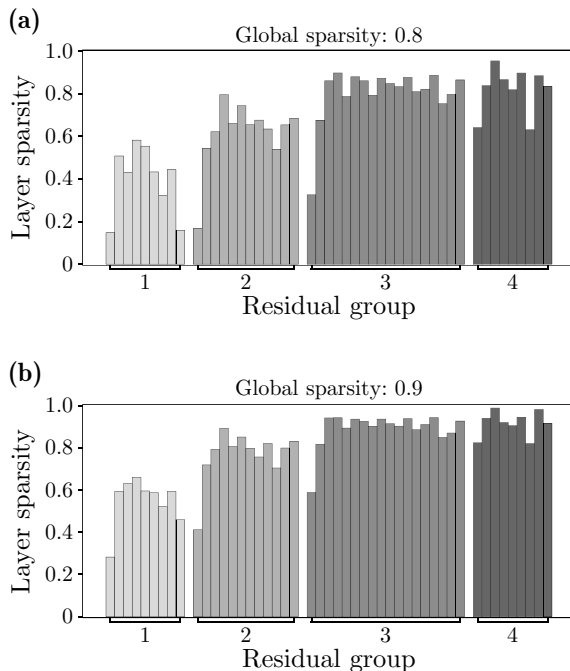
	WRN-28-2 on CIFAR10	Resnet-50 on Imagenet
<i>DeepR</i>	4.466 ± 0.358	5.636 ± 0.218
<i>SET</i>	1.087 ± 0.049	1.009 ± 0.002
<i>Dynamic sparse</i> (ours)	<b>1.083 ± 0.051</b>	<b>1.005 ± 0.004</b>

Median ± standard deviation of wall-clock training epoch times (from 25 epochs) for WRN-28-2 and Resnet-50 for different dynamic reparameterization methods. Results are relative ratios to the epoch time of a sparse network trained without dynamic parameter reallocation. WRN-28-2 is trained on a single, while Resnet-50 on four, Nvidia TITAN Xp GPU(s).



**Figure 1:** WRN-28-2 on CIFAR10. (a) Test accuracy plotted against number of trainable parameters in the sparse models for different methods. Dashed lines are used for the full dense model and for models obtained through compression, whereas all methods that maintain a constant parameter count throughout training and inference are represented by solid lines. Circular symbols mark the median of 5 runs, and error bars are the standard deviation. Parameter counts include all trainable parameters, i.e. parameters in sparse tensors plus all other dense tensors, such as those of batch normalization layers. (b) Breakdown of the final sparsities of the parameter tensors in the three residual blocks that emerged from our dynamic sparse parameterization algorithm (Algorithm 1) at different levels of global sparsity.

post-training pruning of dense networks. A number of recent studies have attempted direct training of sparse networks using *post hoc* information obtained from a pruned model. (Liu et al., 2018) argued that the sparse structure of the pruned model alone suffices to yield high accuracy, i.e. training a model of the same structure, starting with random weights, almost always reaches comparable levels of accuracy as the pruned model. In contrast, (Frankle & Carbin, 2018) argued that a post-compression sparse network structure alone is not sufficient, but necessary, to attain high accuracy, which, as the authors argue, requires both the pruned network connectivity *and* its initial weights when it was trained as part of the dense model pre-compression. In our experiments, we found that neither the *post hoc* sparseness pattern, nor the combination of connectivity and initialization, managed to explain the high performance of sparse

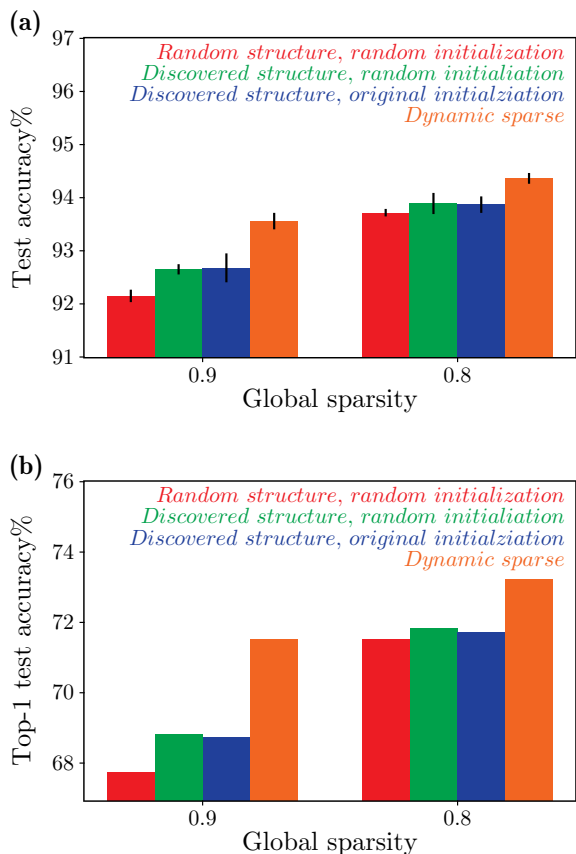


**Figure 2:** layer-wise breakdown of the final parameter tensor sparsities of Resnet-50 trained on Imagenet. (a) At overall sparsity 0.8. (b) At overall sparsity 0.9.

networks produced by our *dynamic sparse* training method. Thus, the value of dynamic parameter reallocation goes beyond discovering trainable sparse network structure: the evolutionary process of structural exploration itself seems helpful for SGD to converge to better weights. Extra work is needed to explain the mechanisms underlying this phenomenon. One hypothesis is that the discontinuous jumps in parameter space when parameters are reallocated across layers helped training escape sharp minima that generalize badly (Keskar et al., 2016).

Structural degrees of freedom are qualitatively different from the degrees of freedom introduced by overparameterization. The latter directly expands the dimensionality of the parameter space in which SGD directly optimizes, whereas structural degrees of freedom are realized and explored using non-differentiable heuristics that only interact indirectly with the dynamics of gradient-based optimization, e.g. regularization pulling weights towards zero causing connections to be pruned. Our results suggest that, for residual CNNs under a given descriptive complexity (i.e. memory storage) constraint, it is better (in the sense of producing models that generalize better) to allocate some memory to describe and explore structural degrees of freedom, than to allocate all memory to conventional weights. This makes a potentially compelling case for hardware acceleration of sparse computations for more parameter efficient training.

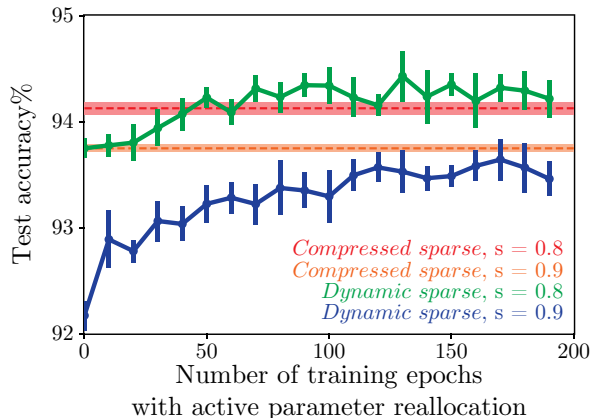
Beside storage (spatial complexity), computational effi-



**Figure 3:** Comparison of training using our dynamic reparameterization method against training a number of related statically parameterized networks. All statically parameterized networks were trained for double the number of epochs used by our method. (a) WRN-28-2 on CIFAR10. Mean and standard deviation from 5 runs. (b) Resnet-50 on Imagenet. Single run.

ciency (temporal complexity) is also of primary concern. Current mainstream computing hardware architectures such as CPUs and GPUs cannot efficiently handle unstructured sparsity patterns. To maintain structured network configurations, various pruning techniques prune a trained model at the level of entire feature maps or layers. Emerging evidence suggests that the resulting pruned networks perform no better than directly-trained thin networks (Liu et al., 2018), calling into question the value of such coarse-grained pruning. We show in Appendix E additional results applying *dynamic sparse* training at an intermediate level of structured sparseness, i.e. pruning  $3 \times 3$  kernel slices. Imposing this sparseness structure led to significantly worse generalization, producing sparse networks performing on par with statically parameterized *thin dense* networks trained for double the number of epochs.

In summary, we show in this paper that it is possible to train deep sparse CNNs directly to reach generalization performances comparable to those achieved by iterative pruning



**Figure 4:** Test accuracies of sparse WRN-28-2 trained on CIFAR10 when dynamic parameter reallocation was switched off at certain epochs. Results are shown for two global sparsity levels: 0.8 and 0.9. Horizontal bands indicate the accuracy of the *compressed sparse* baselines where the band widths represent the standard deviation. For all data points, we ran training for 200 epochs, regardless of when dynamic parameter reallocation was stopped. Mean and standard deviation from 5 independent runs.

and fine-tuning of pre-trained large dense models. The performance level achieved by our proposed method is significantly higher than that achieved by training dense models of the same size. Our method is the first to reallocate free parameters effectively and automatically within and across layers. Furthermore, we show that dynamic exploration of structural degrees of freedom during training is crucial to effective learning. Our work does not contradict the common wisdom that extra degrees of freedom are helpful for training deep networks to achieve better generalization, but it suggests that adding and dynamically exploring structural degrees of freedom is often a more effective and efficient alternative than simply increasing the parameter counts of the model.

## References

- Aimar, A., Mostafa, H., Calabrese, E., Rios-Navarro, A., Tapiador-Morales, R., Lungu, I.-A., Milde, M. B., Corradi, F., Linares-Barranco, A., Liu, S.-C., et al. Nullhop: A flexible convolutional neural network accelerator based on sparse representations of feature maps. *IEEE transactions on neural networks and learning systems*, (99): 1–13, 2018.
- Arjovsky, M., Shah, A., and Bengio, Y. Unitary Evolution Recurrent Neural Networks. nov 2015. URL <http://arxiv.org/abs/1511.06464>.
- Bellec, G., Kappel, D., Maass, W., and Legenstein, R. Deep Rewiring: Training very sparse deep networks. nov 2017. URL <http://arxiv.org/abs/1711.05136>.



- Brutzkus, A., Globerson, A., Malach, E., and Shalev-Shwartz, S. SGD Learns Over-parameterized Networks that Provably Generalize on Linearly Separable Data. oct 2017. URL <http://arxiv.org/abs/1710.10174>.
- Bucilua, C., Caruana, R., and Niculescu-Mizil, A. Model Compression. Technical report, 2006. URL <https://www.cs.cornell.edu/~caruana/compression.kdd06.pdf>.
- Chen, W., Wilson, J. T., Tyree, S., Weinberger, K. Q., and Chen, Y. Compressing Neural Networks with the Hashing Trick. apr 2015. URL <http://arxiv.org/abs/1504.04788>.
- Choromanska, A., Henaff, M., Mathieu, M., Arous, G. B., and LeCun, Y. The Loss Surfaces of Multilayer Networks. nov 2014. URL <http://arxiv.org/abs/1412.0233>.
- Cooper, Y. The loss landscape of overparameterized neural networks. apr 2018. URL <http://arxiv.org/abs/1804.10200>.
- Dai, X., Yin, H., and Jha, N. K. NeST: A Neural Network Synthesis Tool Based on a Grow-and-Prune Paradigm. pp. 1–15, 2017. URL <http://arxiv.org/abs/1711.02017>.
- Dai, X., Yin, H., and Jha, N. K. Grow and Prune Compact, Fast, and Accurate LSTMs. may 2018. URL <http://arxiv.org/abs/1805.11797>.
- Dauphin, Y., Pascanu, R., Gulcehre, C., Cho, K., Ganguli, S., and Bengio, Y. Identifying and attacking the saddle point problem in high-dimensional non-convex optimization. *arXiv*, pp. 1–14, 2014. URL <http://arxiv.org/abs/1406.2572>.
- Denil, M., Shakibi, B., Dinh, L., Ranzato, M., and de Freitas, N. Predicting Parameters in Deep Learning. jun 2013. URL <http://arxiv.org/abs/1306.0543>.
- Dinh, L., Pascanu, R., Bengio, S., and Bengio, Y. Sharp Minima Can Generalize For Deep Nets. 2017. ISSN 1938-7228. URL <http://arxiv.org/abs/1703.04933>.
- Frankle, J. and Carbin, M. The Lottery Ticket Hypothesis: Finding Small, Trainable Neural Networks. mar 2018. URL <http://arxiv.org/abs/1803.03635>.
- Goodfellow, I. J., Vinyals, O., and Saxe, A. M. Qualitatively characterizing neural network optimization problems. dec 2014. URL <http://arxiv.org/abs/1412.6544>.
- Han, S., Mao, H., and Dally, W. J. Deep Compression: Compressing Deep Neural Networks with Pruning, Trained Quantization and Huffman Coding. pp. 1–14, 2015a. doi: [abs/1510.00149/1510.00149](https://doi.org/abs/1510.00149/1510.00149). URL <http://arxiv.org/abs/1510.00149>.
- Han, S., Pool, J., Tran, J., and Dally, W. J. Learning both Weights and Connections for Efficient Neural Networks. jun 2015b. URL <http://arxiv.org/abs/1506.02626>.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep Residual Learning for Image Recognition. *Arxiv.Org*, 7(3):171–180, 2015. ISSN 1664-1078. doi: [10.3389/fpsyg.2013.00124](https://doi.org/10.3389/fpsyg.2013.00124). URL <http://arxiv.org/pdf/1512.03385v1.pdf>.
- Hinton, G., Vinyals, O., and Dean, J. Distilling the Knowledge in a Neural Network. pp. 1–9, 2015. ISSN 0022-2488. doi: [10.1063/1.4931082](https://doi.org/10.1063/1.4931082). URL <http://arxiv.org/abs/1503.02531>.
- Howard, A. G., Zhu, M., Chen, B., Kalenichenko, D., Wang, W., Weyand, T., Andreetto, M., and Adam, H. MobileNets: Efficient Convolutional Neural Networks for Mobile Vision Applications. apr 2017. URL <http://arxiv.org/abs/1704.04861>.
- Huang, Z. and Wang, N. Data-Driven Sparse Structure Selection for Deep Neural Networks. jul 2017. URL <https://arxiv.org/abs/1707.01213>.
- Hubara, I., Courbariaux, M., Soudry, D., El-Yaniv, R., and Bengio, Y. Quantized Neural Networks: Training Neural Networks with Low Precision Weights and Activations. sep 2016. URL <http://arxiv.org/abs/1609.07061>.
- Im, D. J., Tao, M., and Branson, K. An empirical analysis of the optimization of deep network loss surfaces. dec 2016. URL <http://arxiv.org/abs/1612.04010>.
- Jaderberg, M., Vedaldi, A., and Zisserman, A. Speeding up Convolutional Neural Networks with Low Rank Expansions. may 2014. URL <http://arxiv.org/abs/1405.3866>.
- Keskar, N. S., Mudigere, D., Nocedal, J., Smelyanskiy, M., and Tang, P. T. P. On Large-Batch Training for Deep Learning: Generalization Gap and Sharp Minima. sep 2016. URL <http://arxiv.org/abs/1609.04836>.
- Krizhevsky, A., Sutskever, I., and Hinton, G. E. ImageNet Classification with Deep Convolutional Neural Networks. Technical report, 2012.

- Li, C., Farkhoor, H., Liu, R., and Yosinski, J. Measuring the Intrinsic Dimension of Objective Landscapes. apr 2018. URL <http://arxiv.org/abs/1804.08838>.
- Liao, Q. and Poggio, T. Theory of Deep Learning II: Landscape of the Empirical Risk in Deep Learning. *arXiv*, mar 2017. URL <http://arxiv.org/abs/1703.09833>.
- Lin, M., Chen, Q., and Yan, S. Network In Network. *arXiv preprint*, pp. 10, 2013. ISSN 03029743. doi: 10.1109/ASRU.2015.7404828. URL <https://arxiv.org/pdf/1312.4400.pdf><http://arxiv.org/abs/1312.4400>.
- Liu, Z., Sun, M., Zhou, T., Huang, G., and Darrell, T. Rethinking the Value of Network Pruning. oct 2018. URL <http://arxiv.org/abs/1810.05270>.
- Luo, J.-H., Wu, J., and Lin, W. ThiNet: A Filter Level Pruning Method for Deep Neural Network Compression. jul 2017. URL <http://arxiv.org/abs/1707.06342>.
- McDonnell, M. D. Training wide residual networks for deployment using a single bit for each weight. feb 2018. URL <http://arxiv.org/abs/1802.08530>.
- Mocanu, D. C., Mocanu, E., Stone, P., Nguyen, P. H., Gibescu, M., and Liotta, A. Scalable training of artificial neural networks with adaptive sparse connectivity inspired by network science. *Nature Communications*, 9(1):2383, dec 2018. ISSN 2041-1723. doi: 10.1038/s41467-018-04316-3. URL <http://www.nature.com/articles/s41467-018-04316-3>.
- Moczulski, M., Denil, M., Appleyard, J., and de Freitas, N. ACDC: A Structured Efficient Linear Layer. nov 2015. URL <http://arxiv.org/abs/1511.05946>.
- Narang, S., Elsen, E., Diamos, G., and Sengupta, S. Exploring Sparsity in Recurrent Neural Networks. apr 2017. URL <http://arxiv.org/abs/1704.05119>.
- Novak, R., Bahri, Y., Abolafia, D. A., Pennington, J., and Sohl-Dickstein, J. Sensitivity and Generalization in Neural Networks: an Empirical Study. feb 2018. URL <http://arxiv.org/abs/1802.08760>.
- Poggio, T., Kawaguchi, K., Liao, Q., Miranda, B., Rosasco, L., Boix, X., Hidary, J., and Mhaskar, H. Theory of Deep Learning III: explaining the non-overfitting puzzle. 2017. URL <http://arxiv.org/abs/1801.00173>.
- Salimans, T. and Kingma, D. P. Weight Normalization: A Simple Reparameterization to Accelerate Training of Deep Neural Networks. feb 2016. URL <http://arxiv.org/abs/1602.07868>.
- Sifre, L. and Mallat, S. *Rigid-motion scattering for image classification*. PhD thesis, Citeseer, 2014.
- Simonyan, K. and Zisserman, A. Very Deep Convolutional Networks for Large-Scale Image Recognition. sep 2014. URL <http://arxiv.org/abs/1409.1556>.
- Sindhwani, V., Sainath, T. N., and Kumar, S. Structured Transforms for Small-Footprint Deep Learning. oct 2015. URL <http://arxiv.org/abs/1510.01722>.
- Thomas, A. T., Gu, A., Dao, T., Rudra, A., and Christopher, R. Learning invariance with compact transforms. pp. 1–7, 2018.
- Treister, E., Ruthotto, L., Sharni, M., Zafrani, S., and Haber, E. Low-Cost Parameterizations of Deep Convolution Neural Networks. may 2018. URL <http://arxiv.org/abs/1805.07821>.
- Wu, L., Zhu, Z., and E, W. Towards Understanding Generalization of Deep Learning: Perspective of Loss Landscapes. jun 2017. URL <http://arxiv.org/abs/1706.10239>.
- Yang, Z., Moczulski, M., Denil, M., de Freitas, N., Smola, A., Song, L., and Wang, Z. Deep Fried Convnets. dec 2014. URL <http://arxiv.org/abs/1412.7149>.
- Zagoruyko, S. and Komodakis, N. Wide Residual Networks. may 2016. URL <http://arxiv.org/abs/1605.07146>.
- Zhang, C., Bengio, S., Hardt, M., Recht, B., and Vinyals, O. Understanding deep learning requires rethinking generalization. nov 2016. URL <http://arxiv.org/abs/1611.03530>.
- Zhang, C., Liao, Q., Rakhlin, A., Miranda, B., Golowich, N., and Poggio, T. Theory of Deep Learning IIB: Optimization Properties of SGD. jan 2018a. URL <http://arxiv.org/abs/1801.02254>.
- Zhang, T., Ye, S., Zhang, K., Tang, J., Wen, W., Fardad, M., and Wang, Y. A Systematic DNN Weight Pruning Framework using Alternating Direction Method of Multipliers. apr 2018b. URL <http://arxiv.org/abs/1804.03294>.
- Zhu, M. and Gupta, S. To prune, or not to prune: exploring the efficacy of pruning for model compression. 2017. URL <http://arxiv.org/abs/1710.01878>.