

---

# Dense for the Price of Sparse: Improved Performance of Sparsely Initialized Networks via a Subspace Offset

---

Ilan Price<sup>1 2</sup> Jared Tanner<sup>1 2</sup>

## Abstract

That neural networks may be pruned to high sparsities and retain high accuracy is well established. Recent research efforts focus on pruning immediately after initialization so as to allow the computational savings afforded by sparsity to extend to the training process. In this work, we introduce a new ‘DCT plus Sparse’ layer architecture, which maintains information propagation and trainability even with as little as 0.01% trainable kernel parameters remaining. We show that standard training of networks built with these layers, and pruned at initialization, achieves state-of-the-art accuracy for extreme sparsities on a variety of benchmark network architectures and datasets. Moreover, these results are achieved using only simple heuristics to determine the locations of the trainable parameters in the network, and thus without having to initially store or compute with the full, unpruned network, as is required by competing prune-at-initialization algorithms. Switching from standard sparse layers to DCT plus Sparse layers does not increase the storage footprint of a network and incurs only a small additional computational overhead.

## 1. Introduction

It is well established that neural networks can be pruned extensively while retaining high accuracy; see (Blalock et al., 2020; Liu et al., 2020) for recent reviews. Sparse networks have significant potential benefits in terms of the memory and computational costs of training and applying large networks, as well as the cost of communication between servers and edge devices in the context of federated learning. Consequently research on pruning techniques has garnered significant momentum over the last few years.

---

<sup>1</sup>University of Oxford <sup>2</sup>The Alan Turing Institute. Correspondence to: Ilan Price <ilan.price@maths.ox.ac.uk>.

## 1.1. Competing Priorities for Sparse Networks

Traditional pruning algorithms, which prune after or during training, result in a final network with a small storage footprint and fast inference (Gale et al., 2019). However, since these methods initialize networks as dense, and initially train them as such (only to slowly reduce the number of parameters), the overall storage and computational costs of training remain approximately those of a dense network.

For the benefits of sparsity to extend to training, the network must be pruned before training starts. In (Frankle & Carbin, 2019), and many works since (Frankle et al., 2020; Malach et al., 2020), researchers have shown the existence of ‘lottery tickets’ – sparse sub-networks of randomly initialized dense networks, that can be trained on their own from scratch to achieve accuracy similar to that of the full network. This has inspired a surge in recent work on techniques to efficiently prune networks directly at initialization, to identify these trainable, sparse sub-networks.

Research on prune-at-initialization (PaI) methods has progressed rapidly and achieved impressive test accuracy with well below 1% of the network parameters, see (Tanaka et al., 2020; de Jorge et al., 2021) and Section 5. However, almost all current PaI algorithms involve the computation of ‘sensitivity scores’ (or a comparable metric) for all candidate parameters in the dense network, which are then used to decide which parameters to prune. Thus, despite a less computationally demanding training procedure, these methods still require the capacity to store, and compute with, the full network on the relevant device (see Table 1).

Ideally, starting with dense networks which are then pruned would be avoided entirely, and only those parameters to be trained would be initialized. The only method to date that can achieve this is random pruning, since it is equivalent to initializing a sparse network with randomly selected sparse support. For high sparsities, however, random pruning achieves significantly lower accuracy than other methods, for details see Section 5.

Table 1 summarises the performance of the current state-of-the-art pruning algorithms in terms of the various competing priorities for sparse networks: accuracy, storage footprint, and computational complexity.

Table 1. Summary of state-of-the-art prune-at-initialization (PaI) methods in comparison with the DCT plus Sparse (DCTpS) approach presented in this paper. Uniform random pruning is included as baseline, and Iterative Magnitude Pruning (IMP), though not a PaI algorithm, is included for comparison. The table considers only the pruning of network weights, not bias or batchnorm parameters, which is the focus of prior PaI work and of this paper. Let  $N$  denote the total number of parameters in the full networks’ weights,  $P \in (0, 1)$  is the global overall density of the weights tensors, and  $k$  is the number of iterations used in a PaI algorithm. The computational cost in the table refers to the cost of each individual matrix-vector product involved in feedforward and convolutional layers, with flattened weights tensor of size  $m \times n$  with density  $p$ , with  $q = \max(m, n)$  (see Section 5.4 for more details). We report the drop in accuracy relative to the dense baseline at different global sparsities. Quantities which decrease from  $\mathcal{O}(C)$  to  $\mathcal{O}(D)$  during training are denoted by  $\mathcal{O}(C) \xrightarrow{t} \mathcal{O}(D)$ .

	$\Delta$ Accuracy, ResNet50 on CIFAR100			Computational Cost			Network Size on Device		
$P =$	0.01	0.001	0.0001	At init.	Training	Final	At init.	Training	Final
Random	-11.9%	-66%	-66%	0	$\mathcal{O}(pmn)$	$\mathcal{O}(pmn)$	$\mathcal{O}(PN)$	$\mathcal{O}(PN)$	$\mathcal{O}(PN)$
IMP	*	*	*	0	$\mathcal{O}(mn) \xrightarrow{t} \mathcal{O}(pmn)$	$\mathcal{O}(pmn)$	$\mathcal{O}(N)$	$\mathcal{O}(N) \xrightarrow{t} \mathcal{O}(PN)$	$\mathcal{O}(PN)$
FORCE	-6.6%	-26.9%	-62.4%	$\mathcal{O}(mnk)$	$\mathcal{O}(pmn)$	$\mathcal{O}(pmn)$	$\mathcal{O}(N)$	$\mathcal{O}(PN)$	$\mathcal{O}(PN)$
SynFlow	-6.2%	-31.6%	-60.4%	$\mathcal{O}(mnk)$	$\mathcal{O}(pmn)$	$\mathcal{O}(pmn)$	$\mathcal{O}(N)$	$\mathcal{O}(PN)$	$\mathcal{O}(PN)$
<b>Ours</b>	-5.8%	-15%	-22.8%	0	$\mathcal{O}(q \log q + pmn)$	$\mathcal{O}(q \log q + pmn)$	$\mathcal{O}(PN)$	$\mathcal{O}(PN)$	$\mathcal{O}(PN)$

## 1.2. Matching Sparsity vs. Extreme Sparsity

(Frankle et al., 2021) note the distinction between what they call ‘matching sparsities’, at which the resulting pruned networks retain (approximately) the same performance as the full dimensional baseline, and ‘extreme sparsities’, at which there is a trade-off between sparsity and performance. Attention is increasingly being paid to the latter regime, which is especially relevant for resource-constrained settings, in which trade-offs may be necessary or considered worthwhile. A crucial question in the extreme-sparsity setting is the rate of accuracy drop-off as sparsity is increased. Prior algorithms like SNIP (Lee et al., 2019) and GraSP (Wang et al., 2020) display gradual accuracy decrease up to a point, but then reach a sparsity where accuracy rapidly collapses to random guessing. The primary improvements achieved by the most recent algorithms, FORCE (de Jorge et al., 2021) and SynFlow (Tanaka et al., 2020), is to extend that gradual performance degradation to significantly higher sparsity. The DCTpS method proposed here, too, avoids this ‘cliff-like’ drop-off in accuracy, exhibiting an even more gradual decrease in performance at extreme sparsities, resulting in superior performance in this extreme-sparsity regime.

## 1.3. Contributions

In this manuscript we introduce a new neural network layer architecture with which networks can be initialized and trained in an extremely low-dimension parameter space.

These layers are constructed as the sum of a dense *offset* matrix which need not be stored and has a fast transform, plus a sparse matrix of trainable stored parameters, denoted as DpS to abridge ‘Dense plus Sparse’. Consequently, the resulting networks are *in effect* dense, but require the storage of a sparse network with potentially extremely few trainable parameters, and incur the computational cost of very sparse networks. This effective density allows information to continue to propagate through the network at low trainable densities, avoiding unnecessary performance collapse. The aforementioned properties are obtained as follows:

- The neural network layer architectures introduced here are the sum of a discrete cosine transform (DCT) matrix and a sparse matrix, denoted ‘DCT plus Sparse’ (DCTpS). These layers have the same memory footprint as a standard sparse tensor, and a low additional quasi-linear computational overhead above that of sparse layers.
- The sparse trainable matrices from all layers are assigned an equal number of trainable parameters, and within each the support is randomly chosen - avoiding any initial storage of, or computation with the dense network.
- A variety of network architectures using these layers are trained to achieve high accuracy, in particular in the extremely sparse regime with weight-matrix den-

sity as small as  $P = 0.0001$ , where they significantly outperform prior state-of-the-art methods; for example by up to 37% on ResNet50 applied to CIFAR100.

## 2. Prior Prune-at-Initialization (PaI) methods

Neural network pruning has a large and rapidly growing literature; for wider ranging reviews of neural network pruning see (Gale et al., 2019; Blalock et al., 2020; Liu et al., 2020). PaI is the subset of pruning research most directly comparable with the ‘DCT plus Sparse’ networks presented here. For conciseness, we limit our discussion to the most competitive PaI techniques.

The most successful PaI methods determine which entries to prune by computing a synaptic saliency score vector (Tanaka et al., 2020) of the form

$$G(\mathbf{w}) = \frac{\partial \mathcal{R}}{\partial \mathbf{w}} \odot \mathbf{w}, \quad (1)$$

where  $\mathcal{R}$  is a scalar function,  $\mathbf{w}$  is the vector of network parameters, and  $\odot$  denotes the Hadamard product. Those parameters with the lowest scores are pruned.

SNIP (Lee et al., 2019) sets out to prune weights whose removal will minimally affect the training loss at initialization. They suggest ‘connection sensitivity’ as the appropriate metric, with  $G(\mathbf{w}) = |\frac{\partial \mathcal{L}(\mathbf{w})}{\partial \mathbf{w}} \odot \mathbf{w}|$ , where  $\mathcal{L}$  is the training loss.

GraSP (Wang et al., 2020) instead maximises the gradient norm after pruning. The resulting saliency scores for each parameter are calculated via a Taylor expansion of the gradient norm about the dense initialization, resulting in  $G(\mathbf{w}) = -\left(H \frac{\partial \mathcal{L}(\mathbf{w})}{\partial \mathbf{w}}\right) \odot \mathbf{w}$ , where  $H$  is the Hessian of  $\mathcal{L}$ .

FORCE (and a closely related method, iterative SNIP) (de Jorge et al., 2021), like GraSP, take into account the interdependence between parameters so as to predict their importance *after pruning*. They also note, however, that by relying on a Taylor approximation of the gradient norm, GraSP assumes that the pruned network is a small perturbation away from the full network, which is not the case at high sparsities. Instead they propose letting  $G(\mathbf{w}) = |\frac{\partial \mathcal{L}(\bar{\mathbf{w}})}{\partial \mathbf{w}} \odot \mathbf{w}|$ , where  $\bar{\mathbf{w}}$  is the parameter vector  $\mathbf{w}$  after pruning. They then propose FORCE and Iter-SNIP as iterative algorithms to approximately solve for the score vector  $G$  and gradually prune parameters.

SynFlow (Tanaka et al., 2020) makes use of an alternative objective function  $\mathcal{R} = \mathbb{1}^\top \left( \prod_{l=1}^L |\mathbf{w}^{[l]}| \right) \mathbb{1}$ , where  $|\mathbf{w}^{[l]}|$  is the element-wise absolute value of the parameters in the  $l^{\text{th}}$  layer, and  $\mathbb{1}$  is a vector of ones. This allows them to calculate saliency scores  $G$  without using any training data. Like FORCE, their focus extends to extreme sparsities, and their algorithm is designed to avoid layer collapse (pruning whole

layers in their entirety) at the highest possible sparsities. Together, FORCE and SynFlow are the current state-of-the-art for pre-training pruning to extreme sparsities.

However, recent work (Frankle et al., 2021) has shown that given a particular sparsity pattern identified by SNIP, GraSP or SynFlow, one can shuffle the locations of the allotted trainable parameters within each layer, and train the resulting network to matching or even slightly improved accuracy. In other words, they argue, the success of these methods is due to their layer-wise distribution of the number of trainable parameters, rather than the particular location of the trainable parameters within a layer. This somewhat calls into question the role of the proposed saliency metric used to score the importance of each parameter individually. Further understanding of, and heuristics for, these ideal layer-wise parameter allocations would be complementary and directly beneficial to the aforementioned PaI methods as well as ‘DCT plus Sparse’ presented in Section 4.

## 3. Restricting Network Weights to Random Subspaces

Let  $\mathbf{w} \in \mathbb{R}^n$  denote the full vector of network parameters. The number of trainable parameters can be reduced by restricting  $\mathbf{w} \in \mathbb{R}^n$  to a  $k$ -dimensional hyperplane such that

$$\mathbf{w} = \mathbf{d} + U\theta, \quad (2)$$

where  $\mathbf{d}$  is an untrainable offset from the origin,  $U \in \mathbb{R}^{n \times k}$  is a fixed subspace embedding, and  $\theta \in \mathbb{R}^k$  is the vector of  $k \ll n$  trainable parameters. A  $k$ -sparse network ( $\mathbf{w}$  being  $k$ -sparse, with support set  $\mathcal{S}$ ), such as those generated by PaI methods, represents the specific case when  $\mathbf{d} = 0$ , and the subspace embedding  $U \in \mathbb{R}^{n \times k}$  is a matrix with one nonzero per column and at most one nonzero per row, with their locations determined by  $\mathcal{S}$  (we denote this structure for  $U$  as ‘ $k$ -sparse disjoint’). In this sparse  $\mathbf{w}$  setting, identifying ‘Lottery Tickets’ – sparse networks (and their initial parameter values) which can be trained to high accuracy from scratch – can be viewed as identifying the appropriate  $U$  and  $\theta_0$ .

The model (2) was explored in (Li et al., 2018) where they showed that it is possible to *randomly* draw the offset  $\mathbf{d}$  and subspace  $U$ , and to retain accuracy comparable to that of the full  $n$  dimensional network by training only the  $k \ll n$  parameters in  $\theta$ . In their work,  $\mathbf{d}$  is drawn from a traditional, say Gaussian, initialization known to have desirable training properties, and  $U$  has geometry-preserving properties similar to drawing  $U$  uniformly from the Grassmannian; for details see (Li et al., 2018) Appendix S7. The smallest possible dimension  $k$  for which such subspace training achieved 90% of the accuracy of a dense network was termed the ‘intrinsic dimension’ of the loss surface, as the ability to successfully train a network in a randomly chosen low dimensional sub-

space indicates some low-dimensional structure in the loss landscape.

In Figure 1, we repeat one of the experiments from (Li et al., 2018), comparing their method, which we denote as ‘Hyperplane Projection’, with random pruning and the aforementioned PaI methods, on Lenet-5 with CIFAR10. The performance of Li et al.’s method stands in stark contrast with the performance of random pruning at initialization, which corresponds to (2) with  $d = 0$  and  $U$  being  $k$ -sparse disjoint, with  $S$  selected uniformly at random. Despite both methods involving training in randomly selected subspaces, ‘Hyperplane Projection’ far outperforms random pruning at  $k \ll n$ . Furthermore, in this low-dimensional regime, ‘Hyperplane Projection’ even outperforms state-of-the-art PaI algorithms.




Figure 1. Different subspace selection methods applied to Lenet-5 (LeCun et al., 1998) and trained on CIFAR10. We report maximum validation accuracy, averaged over three runs, at each subspace dimension (number of trainable parameters).

However, in the context of PaI algorithms it is important to note that despite having the same number of trainable parameters  $k$ , the networks based on the ‘Hyperplane Projection’ model (2), with  $d \in \mathbb{R}^n$  dense, do not afford any memory and computational benefits over a dense network.

In the following section we propose an alternative subspace model to (2), ‘DCT plus Sparse’, which combines the benefit of the dense nonzero offset  $d$  of (Li et al., 2018) with the specially structured sparse  $U$  as used in PaI methods, without needing to store the offset  $d$ . Moreover, we show that state-of-the-art test accuracy is obtained even while selecting the location of the 1-sparse rows in  $U$  according to a simple random equal per layer heuristic, avoiding the initial calculation of parameter saliency scores.

#### 4. DCT plus Sparse (DCTpS) Network Layers

The network parameters  $w$  in (2) which PaI methods sparsify are typically only the network weight matrices as they

usually comprise the largest number of trainable parameters<sup>1</sup>. In the specific case we consider, combining  $d$  dense and non-trainable, with  $U$  being  $k$ -sparse disjoint, the associated weight matrices  $W \in \mathbb{R}^{m \times n}$  comprising  $w$  in (2) can be expressed as  $W = D + S$ , where  $D$  is dense, but fixed (i.e. non-trainable  $d$ ), and  $S$  is sparse, with fixed sparse support (i.e. non-trainable  $U$ ) and trainable values within that support (corresponding to the trainable  $\theta$ ). As  $D$  is dense, the sparse matrix  $S$  can be initialized as zero, and the training of  $S$  corresponds to adjusting only  $k \ll n$  entries within  $D$ . To allow for an additional bulk scaling of  $D$  by a trainable parameter<sup>2</sup>  $\alpha$ , similar to batch normalization (Ioffe & Szegedy, 2015), we consider  $W = \alpha D + S$ .

In order to maintain the low network size on device and to reduce the computational burden of applying  $W$  with a dense component (and  $W^\top$  in the backward pass), we treat the dense offset  $D$  as the action of the discrete cosine transform (DCT) matrix<sup>3</sup> resulting in

$$Wx = \alpha \text{DCT}(x) + Sx, \quad (3)$$

$$W^\top v = \alpha \text{DCT}^{-1}(v) + S^\top v. \quad (4)$$

The DCT can be applied in near linear,  $q \log q$ , computational cost (where  $q = \max(m, n)$ ), and need not be directly stored. Consequently, this layer architecture (3) retains the benefit of  $W$  being dense, while having the on device storage footprint of a sparse network and at the minimal overhead of requiring an additional computational  $q \log q$  cost. We refer to layers parameterized by (3) as ‘DCT plus Sparse’ (DCTpS) layers. There are, of course, many other candidate matrices for  $D$  with the same or similar properties, which together constitute a more general ‘Dense plus Sparse’ (DpS) layer class, but we restrict our attention to DCTpS in this paper, deferring alternative choices of fast transforms to later investigation.

The framework (3) applies equally to convolutional layers. Each step in a convolution can be cast as a matrix-vector product  $Wx$ , where  $x \in \mathbb{R}^{\text{patch.dim}}$  is a vectorised ‘patch’ of the layer input, and  $W \in \mathbb{R}^{\text{out.channels} \times \text{patch.dim}}$  has the filters as its rows. Back-propagation through a 2D convolutional layer involves convolutions with (rotated versions of) the layer’s filters, each step of which can be implemented as  $W\hat{v}$ , where  $\hat{v}$  is a permuted version of a patch  $v$ . See the Supplementary Material for more details.

The state-of-the-art PaI algorithms, such as FORCE and Syn-

<sup>1</sup>Tanaka et al. briefly extend their analysis to batchnorm layers in their Appendix (Section 10) (Tanaka et al., 2020).

<sup>2</sup>We note that the inclusion of an  $\alpha$  scaling parameter for  $D$  is a departure from a standard subspace training model since it enables the re-scaling of different sections of  $d$  independently during training, but it adds expressive power to the network with almost no overhead.

<sup>3</sup>If the input dimension is less or greater than the output dimension, we zero-pad the input or truncate the output respectively.

Flow, require initially storing and computing with densely initialized weight matrices so as to compute saliency scores for each parameter. To avoid these costs, we use only a simple heuristic to determine the locations of the trainable parameters in DCTpS networks: we allocate an equal number of trainable parameters to each layer, and draw the locations of those trainable parameters within each layer uniformly at random.

This ‘Equal per layer’ (EPL) heuristic achieves the basic goal of maintaining some amount of trainability within each layer, but is otherwise naive. While we will show that even something as simple as EPL is sufficient for state-of-the-art results with DCTpS networks, there is likely scope for improved heuristics for the allocation of trainable parameters, which - as noted in (Frankle et al., 2021) - may be the most relevant feature of a PaI method, and thus may further improve performance. In the Supplementary Material, we include experiments with another naive heuristic which distributes parameters evenly across filters, rather than layers, and achieve similar results.

## 5. Experiments

SynFlow (Tanaka et al., 2020) and FORCE (de Jorge et al., 2021) have very recently emerged as the state-of-the-art PaI algorithms in the extreme-sparsity regime, significantly outperforming the prior state-of-the-art methods SNIP and GraSP. We thus focus our experiments on comparing DCTpS with SynFlow and FORCE.<sup>4</sup> A full description of all experimental setups and hyperparameters is included in the Supplementary Material. For all plots in this section, solid lines represent test accuracy averaged over three runs, and shaded regions (though often too small to make out) represent the standard deviation. The dashed black lines denote the unpruned dense network baseline accuracy, while dashed colored lines indicate where an algorithm breaks down and is thus unable to prune the network to the specified sparsity.

**Random Pruning Comparisons:** As noted above, the support sets of trainable nonzero parameters in ‘DCTpS’ networks are selected without any calculations involving the full network. Relevant comparisons in this respect are thus variants of random pruning, since initializing sparse matrices is equivalent to initializing them as dense and pruning randomly. Globally uniform random pruning, which we denote as Sparse (uniform) in Figures 1 - 4; is often included as a baseline in works such as (de Jorge et al., 2021; Lee et al., 2019; Wang et al., 2020)). We include an additional random sparse initialization, Sparse (EPL), which uses the

<sup>4</sup>In a recently published revised version, (de Jorge et al., 2021) includes a closely related variant of FORCE, called Iter-SNIP. The results for both methods are very similar, however, and it suffices to compare our method to FORCE.

same heuristic for distributing trainable parameters as we use for DCTpS networks (with the difference being that the trainable entries in the sparse weight matrices are not initialized as 0, but according to a standard initialization scheme). We note that this Sparse (EPL) heuristic significantly outperforms Sparse (uniform) at higher sparsities, and even matches state-of-the-art pruning methods for densities greater than approximately 1%, see e.g. Figure 2.

### 5.1. Lenet-5

We first consider the small Lenet-5 architecture on CIFAR10 in order to compare the sparse network methods against the ‘Hyperplane Projection’ method of (Li et al., 2018), which is computationally demanding despite having few trainable parameters, due to the nature of its chosen random subspace. Figure 1 illustrates that Hyperplane Projection achieves validation accuracy superior to all PaI methods tested except DCTpS which matches or exceeds its accuracy. The efficacy of the Hyperplane Projection method helps illustrate the value of the affine offset in (2), while the even superior accuracy of DCTpS demonstrates that the offset can be deterministic and the hyperplane sparse and axis-aligned as in (3).

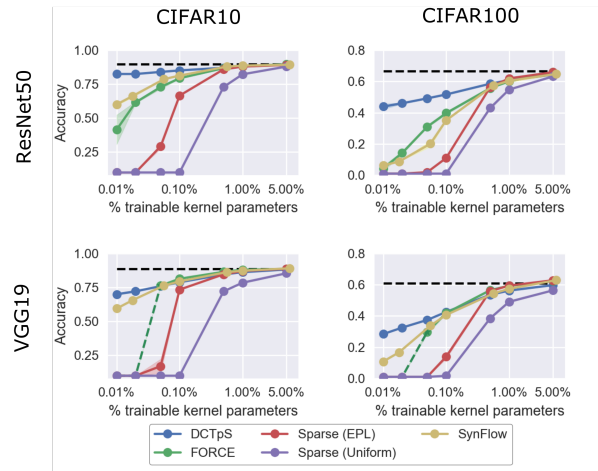


Figure 2. Test accuracy on CIFAR10 and CIFAR100 datasets using sparse ResNet50 and VGG19 architectures. DCT plus Sparse (DCTpS) networks (with EPL parameter allocation) as compared with FORCE, SynFlow, and random pruning methods.

### 5.2. ResNet 50 and VGG19 applied to CIFAR10 and CIFAR100

ResNet50 and VGG19 are selected as the primary architectures to benchmark the PaI methods considered here; this



follows (de Jorge et al., 2021) and allows direct comparison with related experiments conducted therein.

Figure 2 displays the validation accuracy of these architectures, applied to CIFAR10 and CIFAR100 datasets, as a function of the percentage of trainable parameters within the weight matrices determined by the aforementioned PaI algorithms. At 5% density all PaI algorithms are able to obtain test accuracy approximately equal to that of a dense network. Sparse (uniform) and sparse (EPL) initializations exhibit a collapse or significant drop in accuracy once the density drops below 1% and 0.5% respectively, but at greater densities they roughly match or even outperform other methods. It is therefore at densities of  $< 0.5\%$  where the more complex PaI methods become necessary, and in this regime DCTpS has superior or equal test accuracy compared to both SynFlow<sup>5</sup> and FORCE. The superior accuracy of DCTpS is most pronounced as the density decreases to 0.01%; for example, in the case of ResNet applied to CIFAR10, DCTpS achieves accuracy 22% above the next most effective method SynFlow, and moreover retains accuracy in excess of 80%.

The test accuracy of all methods are somewhat lower for VGG19 applied to CIFAR10 and CIFAR100. Up to approximately 0.05% density, each of SynFlow, FORCE, and DCTpS achieve accuracy within approximately 2% of each other. FORCE fails to generate a network for densities below 0.05%, denoted by dashed green lines, while the benefits of DCTpS over SynFlow become apparent for these smaller densities where its test accuracy is approximately 10% and 20% greater on CIFAR10 and CIFAR100 respectively.

One important point to underscore is that in these experiments, as is typical in prior PaI experiments (de Jorge et al., 2021), we only consider the proportion of remaining trainable *weights* in linear and convolutional layers. Large networks like ResNet50 and VGG19 also have bias and batchnorm parameters (making up 0.22% and 0.06% of trainable parameters of the respective architectures). That these are not prunable in our experiments imposes a floor on the overall number of trainable parameters remaining in the network. The plateau in performance exhibited by DCTpS networks in Figure 2 is thus testament to their ability to preserve information flow through the network despite extremely few trainable weights, thereby preserving the capacity of the network endowed by other remaining trainable parameters.

Shrinking the storage footprints of batchnorm networks to their most extreme limits will require the development and incorporation of methods to prune batchnorm parameters

<sup>5</sup>We note that one feature of SynFlow’s saliency scores is that they grow very large for large networks, and thus in order to successfully prune ResNet50 with SynFlow it is necessary to switch from float32 to float64 to avoid overflow.

as well. Such methods can be incorporated into DCTpS networks (and other PaI methods).

### 5.3. MobileNetV2 and Fixup-ResNet110

Next, DCTpS is compared to other PaI algorithms on two architectures which are less overparameterized than ResNet50 and VGG19. First, we consider MobileNetV2 (Sandler et al., 2018), originally proposed as a PaI test-case in (de Jorge et al., 2021), which has approximately 10% of the parameters in ResNet50. Figure 3 (left) shows test accuracy for MobileNetV2 applied to CIFAR10, demonstrating trends similar to ResNet50 and VGG19 in Figure 2, with DCTpS exhibiting superior performance as sparsity increases, retaining approximately 70% accuracy with as few as 0.01% of the networks weights.

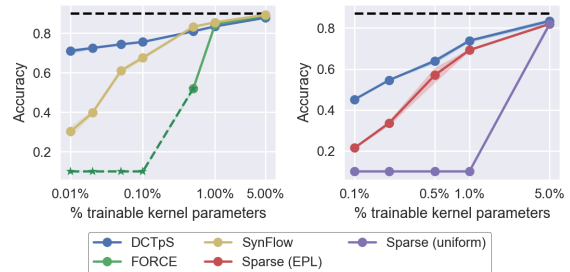


Figure 3. Left: Comparison of DCTpS against FORCE and SynFlow on CIFAR10 with MobileNetV2 (Sandler et al., 2018). Right: Comparison of DCTpS with Random Pruning methods on CIFAR10 with FixupResNet110 (Zhang et al., 2019).

Second, we include experiments for Fixup-ResNet110 (Zhang et al., 2019) on CIFAR10. ‘Fixup’ ResNets were developed in order to enable the efficient training of very deep residual networks without batchnorm, to the same accuracy as similarly sized batchnorm networks. In Fixup-ResNet110, all but 282 of its 1720138 parameters are ‘prunable’, practically eliminating the overall density floor caused by batchnorm parameters in the other large networks considered in this section.

The Fixup initialization involves initializing some layers to zero (in particular the classification layer and the final layer in each residual block), as well as re-scaling the weight layers inside the residual branches. Since our layers are initialized as DCTs, we mimic these effects by setting the  $\alpha$  parameter to 0 or to the appropriate scaling factor. We use the code provided by the authors<sup>6</sup> to obtain the baseline and sparse network results, and create DCTpS Fixup ResNets by simply replacing the Linear and Convolutional layers with their corresponding DCTpS variants, and initializing

<sup>6</sup><https://github.com/hongyi-zhang/Fixup>




Figure 4. Spectrum of the Jacobian of Resnet50 on CIFAR10, pruned to different sparsities with varies methods, at initialization. If a curve does not appear in a given plot, it means that the spectrum was identically zero for that density. The DCTpS plot shows only one curve since its Jacobian does not depend on the number of trainable weights.

appropriately.

Figure 3 (right) illustrates the validation accuracy of DCTpS as well as Sparse (EPL) and Sparse (uniform) initializations. FORCE and SynFlow cannot be directly applied to FixupResNet110 as they both assign a saliency score of 0 to all parameters and consequently have no basis on which to select which entries to prune. At 0.1% density, with only 2000 trainable parameters in total, spread across 110 layers, DCTpS retains approximately 45% test accuracy, outperforming the Sparse (EPL) by more than 20%.

#### 5.4. Run Times and Theoretical Complexity

The operations of both fully connected and convolutional layers can be framed in terms of matrix multiplication, and thus we may discuss the complexity of a DCTpS layer as compared with a standard sparse layer by considering a matrix-vector product  $Wx$ , where  $W \in \mathbb{R}^{m \times n}$ . If  $Wx$  is parameterized as in Equation 3, with  $S$  containing  $pmn$  non-zeros, then theoretically the storage cost of  $W$  is just the storage cost of  $S$ ,  $\mathcal{O}(pmn)$ , while the computational cost of applying the layer (computing  $Wx$ ) becomes  $\mathcal{O}(q \log q + pmn)$ , where  $q = \max(m, n)$ . We note here that while storage requirements can decrease to arbitrarily low levels, depending on  $p$ , once  $p < \mathcal{O}(\log q/q)$ , further sparsification results in only minimal computational savings.

However, current implementations of deep learning packages render these computational and storage gains purely theoretical, for now at least, and thus plots showing run-times and storage costs corresponding to the above expressions are not included. Firstly, most standard deep learning libraries are not optimised for sparse tensor operations, which affects the realisation of the potential benefits of all PaI techniques, as well as our DCTpS approach. Secondly, efficient DCTpS networks would require optimising the fast transforms in these packages, and appropriately building in their auto-differentiation.

## 6. Spectral Analysis and Distribution of Nonzeros

Figure 2 shows that the accuracy of Sparse (uniform) collapses to random guessing as the density decreases from 0.5% to 0.1%, and for Sparse (EPL) this occurs at 0.02%. In both cases the percentage of trainable parameters at which the pruned network becomes un-trainable are those at which the spectrum of their Jacobian ( $J_{ij} = \frac{\partial f_i}{\partial x_j}$ ) becomes equal to 0 at initialization. This can be seen in Figure 4, which displays the associated leading singular values of the Jacobian in the exemplar case of ResNet50 applied to CIFAR10, for different PaI methods at varying densities. As previously mentioned, FORCE and SynFlow do not reduce to random guessing at any sparsity tested with ResNet50, and correspondingly there is no sparsity at which the Jacobian spectrum fully collapses to zero. It appears that the primary factor for trainability of PaI networks is determined by whether the spectrum is or is not 0, as opposed to the scale of the spectrum – SynFlow remains competitive with FORCE even at densities for which SynFlow’s Jacobian has a largest singular value of approximately  $10^{-13}$  whereas the corresponding singular value of FORCE is more than  $10^{10}$  greater.

Since DCTpS networks are always, in effect, dense networks, including at initialization, the spectrum of the Jacobian, shown in Figure 4, does not depend on the sparsity of its trainable weights, and thus does not collapse even as the number of trainable parameters approaches zero. This likely relates to their ability to be trained even with extremely sparse and randomly distributed trainable weights.

It has been alternatively conjectured that these noted collapses in accuracy occur due to ‘layer collapse’, where one or more layers have *all* of their parameters set to zero, as was noted in (Tanaka et al., 2020) for SNIP and GraSP. However, in residual networks, due to the existence of multiple branches through which information may flow, it is possible




Figure 5. Number of remaining parameters (y-axis) in the prunable layers (x-axis) in ResNet50 after pruning with different PaI methods at a variety of densities. Each color curve represents a different global density. The insets zoom in on portions of the plot with the lowest totals, to illustrate where methods do or do not prune weight tensors in their entirety.

to prune multiple weights tensors in their entirety without collapsing performance. Indeed this phenomenon is observed in Figure 5, which includes plots of the total number of trainable parameters per layer determined by different PaI methods at different sparsities, applied to ResNet50 on CIFAR10. Comparing Figure 5 to the test accuracy of the corresponding experiments in Figure 2, shows that pruning all parameters in one or more weight tensors is neither necessary nor sufficient for accuracy collapse. Neither SynFlow nor FORCE exhibit complete performance collapse at any density, despite both methods pruning multiple layers completely. Conversely, Figure 2 shows Sparse (EPL) on ResNet50 reduces to random guessing at a density of

0.02% or less, despite the fact that each single layer contains approximately 100 trainable parameters.

The number and distribution of nonzeros generated by PaI methods has recently been investigated in depth by (Frankle et al., 2021); and in particular, the value of carefully selecting sparse trainable parameters with PaI methods. It was observed that given a particular sparsity pattern identified by PaI methods, the location of the nonzeros within a layer can be shuffled and the resulting network can be trained to similar or even improved accuracy, suggesting the success of a PaI algorithm may be determined primarily by the distribution of the number of nonzeros per layer rather than which entries within the layer are selected. Figure 5 illustrates that SynFlow, for densities down to 0.1%, allocates trainable parameters approximately equally per layer (with 4 notable exceptions, which turn out to correspond to those shortcut connections which were prunable). Yet despite this rough similarity in distribution at 0.1% density, we observe substantially different test accuracy for SynFlow and Sparse (EPL).

## 7. Conclusions and Further Work

In this work we have shown that adding a layer-wise offset to sparse subspace training significantly improves validation accuracy in the extreme-sparsity regime. DCT plus Sparse (DCTpS) layers provide an elegant way of achieving this offset with no extra storage cost, and only a small computational overhead. Moreover we have shown that simple heuristics can be used to randomly select the support sets for their sparse trainable weight tensors, avoiding any initial storage of, or computation with the full network. DCTpS networks achieve state-of-the-art results at extreme trainable sparsities, and are competitive with the state-of-the-art at lower sparsities.

There are numerous clear avenues for extending and complementing this research. As noted, research on better heuristics or other ways to choose the trainable sparse support in the DCTpS layers may improve performance beyond our simple EPL heuristic. Moreover, this should be combined with research on pruning or removing batchnorm parameters, since DCTpS layers only enable the pruning of trainable weight tensors. Furthermore, there may well be an optimal initialization of the  $\alpha$  parameter in DCTpS layers, and it may not need to be trained. As the scaling parameter of the network weight tensors, this research would be analogous to work on the optimal scale parameter to use when initializing Gaussian weights, of which there is plenty (Glorot & Bengio, 2010; He et al., 2015; Xiao et al., 2018). Finally, the use of other fast deterministic transforms, or other even more efficient ways to implement a dense offset, may yield further improvements.



---

## References

- Blalock, D., Ortiz, J. J. G., Frankle, J., and Gutttag, J. What is the state of neural network pruning? *arXiv preprint arXiv:2003.03033*, 2020.
- Boué, L. Deep learning for pedestrians: backpropagation in cnns. *arXiv preprint arXiv:1811.11987*, 2018.
- de Jorge, P., Sanyal, A., Behl, H., Torr, P., Rogez, G., and Dokania, P. K. Progressive skeletonization: Trimming more fat from a network at initialization. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=9GsFOUyUPi>.
- Frankle, J. and Carbin, M. The lottery ticket hypothesis: Finding sparse, trainable neural networks. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=rJl-b3RcF7>.
- Frankle, J., Dziugaite, G. K., Roy, D., and Carbin, M. Linear mode connectivity and the lottery ticket hypothesis. In *International Conference on Machine Learning*, pp. 3259–3269. PMLR, 2020.
- Frankle, J., Dziugaite, G. K., Roy, D., and Carbin, M. Pruning neural networks at initialization: Why are we missing the mark? In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=Ig-VyQc-MLK>.
- Gale, T., Elsen, E., and Hooker, S. The state of sparsity in deep neural networks. *arXiv preprint arXiv:1902.09574*, 2019.
- Glorot, X. and Bengio, Y. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pp. 249–256. JMLR Workshop and Conference Proceedings, 2010.
- Hadjis, S., Abuzaid, F., Zhang, C., and Ré, C. Caffe con troll: Shallow ideas to speed up deep learning. In *Proceedings of the Fourth Workshop on Data analytics in the Cloud*, pp. 1–4, 2015.
- He, K., Zhang, X., Ren, S., and Sun, J. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE international conference on computer vision*, pp. 1026–1034, 2015.
- Ioffe, S. and Szegedy, C. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International conference on machine learning*, pp. 448–456. PMLR, 2015.
- Krizhevsky, A., Hinton, G., et al. Learning multiple layers of features from tiny images. 2009.
- LeCun, Y., Bottou, L., Bengio, Y., and Haffner, P. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- Lee, N., Ajanthan, T., and Torr, P. Snip: Single-shot network pruning based on connection sensitivity. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=BlVZqjAcYX>.
- Li, C., Farkhor, H., Liu, R., and Yosinski, J. Measuring the intrinsic dimension of objective landscapes. In *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=ryup8-WCW>.
- Liu, J., Tripathi, S., Kurup, U., and Shah, M. Pruning algorithms to accelerate convolutional neural networks for edge applications: A survey. *arXiv preprint arXiv:2005.04275*, 2020.
- Malach, E., Yehudai, G., Shalev-Schwartz, S., and Shamir, O. Proving the lottery ticket hypothesis: Pruning is all you need. In *International Conference on Machine Learning*, pp. 6682–6691. PMLR, 2020.
- Sandler, M., Howard, A., Zhu, M., Zhmoginov, A., and Chen, L.-C. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 4510–4520, 2018.
- Tanaka, H., Kunin, D., Yamins, D. L., and Ganguli, S. Pruning neural networks without any data by iteratively conserving synaptic flow. *Advances in Neural Information Processing Systems*, 33, 2020.
- Wang, C., Zhang, G., and Grosse, R. Picking winning tickets before training by preserving gradient flow. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=SkgsACVKPH>.
- Wu, J., Zhang, Q., and Xu, G. Tiny imagenet challenge, 2017.
- Xiao, L., Bahri, Y., Sohl-Dickstein, J., Schoenholz, S., and Pennington, J. Dynamical isometry and a mean field theory of cnns: How to train 10,000-layer vanilla convolutional neural networks. In *International Conference on Machine Learning*, pp. 5393–5402. PMLR, 2018.
- Zhang, H., Dauphin, Y. N., and Ma, T. Residual learning without normalization via better initialization. In *International Conference on Learning Representations*,

---

2019. URL <https://openreview.net/forum?id=H1gsz30cKX>.

## A. Additional Experiments

### A.1. Tiny Imagenet

Figure 6 compares the performance of DCTpS with SynFlow and FORCE<sup>7</sup> on Tiny Imagenet with ResNet18. DCTpS obtains higher validation accuracy than both FORCE and SynFlow for all densities less than or equal to 1%, and by 0.1% density DCTpS is outperforming them by approximately 10% accuracy. This confirms that the superior accuracy of DCTpS networks at low densities is maintained when the difficulty of the problem is scaled up.




Figure 6. Comparing DCTpS with FORCE and SynFlow on Tiny Imagenet with ResNet18.

### A.2. Equal-per-filter (EPF) support distribution

All DCTpS experiments in Section 5 used the EPL heuristic to distribute trainable parameters between layers. Another naive heuristic which achieves the basic goal of maintaining some trainability throughout the network is an ‘Equal per Filter’ (EPF) approach: given a specified sparsity, the total number of trainable parameters for the whole network is calculated, and divided equally between all convolutional filters (or rows in the case of linear layers). Within each filter, the locations of those trainable parameters is chosen uniformly at random.

Figure 7 highlights that the EPL and EPF heuristics result in substantially and qualitatively different layer-wise parameter allocations. Nevertheless, Figure 8 shows that both methods achieve very similar accuracy across all tested densities, though EPF consistently performs marginally worse. This observation lends further support to the hypothesis that, provided a suitable offset  $d$  is used, there is a large class of

<sup>7</sup>We note that FORCE was particularly unstable in these experiments, failing to prune the network to the required density in at least one of its three runs, at every density less than 0.5%. In these cases, accuracy is averaged over only those runs in which FORCE succeeded. At 0.01% density, FORCE failed on all three runs.

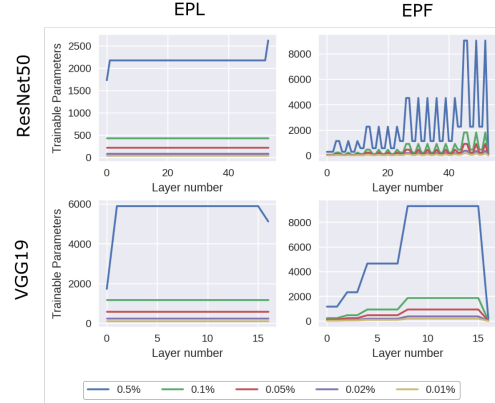


Figure 7. The total number of trainable parameters per prunable layer determined by EPL and EPF heuristics, on ResNet50 and VGG19 with 10 output classes.

subspace embeddings from which it suffices to draw  $U$  randomly to achieve high accuracy, and in particular that this class includes  $k$ -sparse disjoint  $U$  with a variety of support distributions.

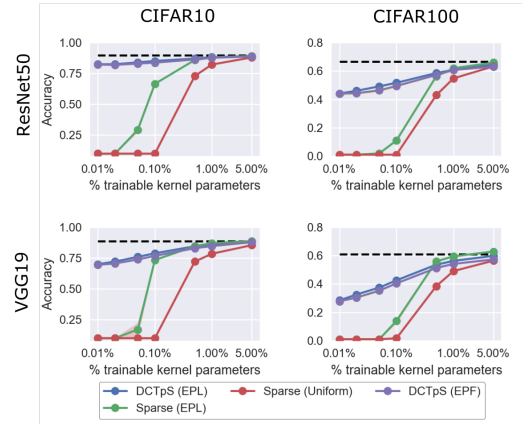


Figure 8. Comparative accuracy of EPF and EPL heuristics for DCTpS networks.

### A.3. Training with SGD as opposed to Adam

The best test accuracy for large networks like ResNet and VGG is typically obtained by using SGD with momentum and a specified learning rate schedule, rather than adaptive methods like Adam. However, the results obtained with SGD are sensitive to hyperparameters like the initial learning rate and the learning rate schedule. Adam, though it tends to achieve lower final accuracy, is less sensitive to

these hyperparameters. This makes Adam a sensible training algorithm for experiments in which the goal is to compare the relative drop in accuracy caused by one or other pruning method, as opposed to a goal of achieving maximum possible accuracy overall. Thus Adam is used, with default settings, as the training algorithm for the experiments presented in Section 5. In order to preserve comparison with prior work, and to illustrate that our results are not unique to Adam, additional experiments using SGD with momentum on ResNet50, VGG19, MobileNetV2 and FixUpResNet110 are included here. A single, coarse sweep of base learning rates [0.1, 0.07, 0.05, 0.03, 0.01] was done with a DCTpS ResNet50 applied to CIFAR100, at 1% density, to select a base learning rate of 0.03, which was then used to train all DCTpS architectures, at all densities, without further fine-tuning. For PaI on standard architectures, a base learning rate of 0.1 was used as done in prior work (de Jorge et al., 2021).

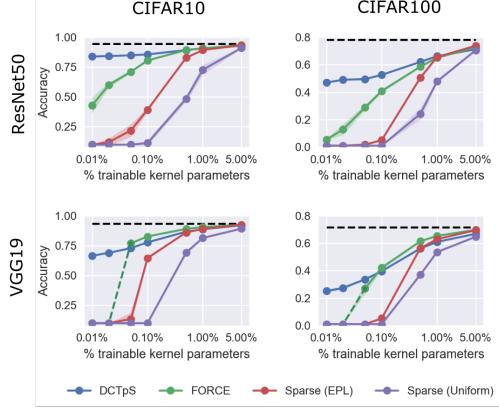


Figure 9. Experiments training ResNet50 and VGG19 with SGD (with momentum), on CIFAR10 and CIFAR100.

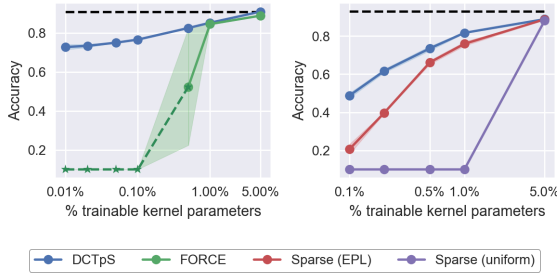


Figure 10. Experiments training MobileNetV2 and FixupResNet110 with SGD (with momentum), on CIFAR10.

Test accuracy is shown in Figure 9 for ResNet50 and VGG19, and Figure 10 for MobileNetV2 and FixupResNet110<sup>8</sup>. The results exhibit qualitatively similar trends

<sup>8</sup>SynFlow was not able to be included in these additional supple-

to those observed in Section 5’s Figures 2 and 3, though with slightly higher overall accuracy, in particular at higher densities, as expected.

## B. DCTpS Implementation

In the code used to run the experiments in this paper, the DCT components of DCTpS layers have been implemented by setting the tensor  $W$  to be the DCT matrix (matrix of DCT basis vectors), as opposed to implementing them via fast transforms. This is because deep learning libraries, as currently implemented, are optimised for the former rather than the latter.

**Linear Layers:** To be precise, in Linear layers with input  $\in \mathbb{R}^n$  and output  $\in \mathbb{R}^m$ , with  $q = \max(m, n)$ , the DCT matrix  $W \in \mathbb{R}^{q \times q}$  is formed and then truncated to size  $m \times n$  by removing the surplus right-most columns (if  $m > n$ ) or bottom-most rows (if  $m < n$ ). Multiplication by this matrix is equivalent to a DCT, with a zero-padded input if  $m > n$ , or a truncated output if  $m < n$ .

**Convolutional Layers:** In a convolutional layer, with  $m$  ( $k \times k \times n$ ) filters,  $q = \max(k^2n, m)$ , the DCT matrix  $W \in \mathbb{R}^{q \times q}$  is formed and truncated to size  $k^2n \times m$  and reshape appropriately. In the accompanying code, a simple test script is provided to confirm that our implementation indeed computes the DCT of each patch.

Figure 11 provides a simple visualisation of how this is compatible with the efficiency of DCTpS layers, if implemented correctly. In the forward pass, each step of the convolution involves taking a patch of the image, and - for each filter - computing the sum of the elementwise product of the patch and the filter. Flattening (commonly known as ‘lowering’) the filters and the input patch, this is equivalent to a matrix-vector product (and indeed convolutional layers are commonly implemented with this ‘lower  $\rightarrow$  matrix multiply  $\rightarrow$  lift’ approach (Hadjis et al., 2015)). The DCT part of DCTpS convolutional layers set this flattened matrix of filters to be the DCT matrix, and is thus equivalent to computing the DCT of each patch.

This applies equally for the backward-pass, since backpropagation through convolutional layers involves convolutions as well. Let  $h$  be the layer input,  $y$  be the output of a layer with a single  $2 \times 2$  filter  $F$ , and let  $L$  be the loss. Calculating  $\frac{\partial L}{\partial h}$  involves calculating  $\frac{\partial L}{\partial y} * \text{Rot}_{180}(F)$ . Again, each step in this convolution is equivalent to an inner product between the original filter, and a flattened, *permuted* patch of  $\frac{\partial L}{\partial y}$ , see Figure 12. This generalises for larger filters (Bou  , 2018).

mentary experiments having only recently been published (Tanaka et al., 2020).




Figure 11. Illustration of the matrix-vector product involved in each step of a 2D convolutional layer. Computing the DCT of each patch is equivalent to taking  $W$  to be a DCT matrix.




Figure 12. Back-propagation involves inner products with a layer's filters. In this figure,  $d$  represents a patch of  $\frac{\partial L}{\partial y}$ , and  $F$  is one of the layers convolutional filters.

### C. Parameter Allocation by SynFlow

It was noted in Section 6 that when applied to ResNet50 for CIFAR10 at modest sparsities, SynFlow pruned fully those residual connections which were prunable (those implemented as trainable,  $1 \times 1$  convolutions), and in the remaining layers it appeared to approximate the EPL heuristic. Figure 13 shows that this observation also applies to other architectures with different numbers of output classes. It is

particularly striking the extent to which SynFlow applied to VGG19 leaves unpruned an approximately equal number of parameters per layer. On ResNet18, the pruning is observed to have the same structure per layer as on ResNet50 - a roughly equal number of trainable parameters per layer, except for the prunable shortcut connections, which are pruned completely. This behaviour is also present to some extent on MobileNetV2, though with larger oscillations around a central value, and a breakdown of this behaviour at lower densities, at which point multiple layers begin to be pruned in their entirety.

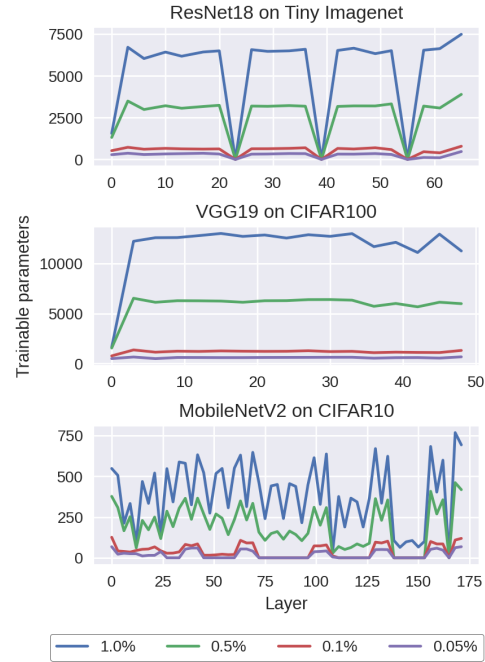


Figure 13. Total number of trainable parameters per layer, after pruning at initialization with SynFlow.

### D. Cases where SynFlow and FORCE cannot be applied

#### D.1. Extreme Sparsities (FORCE)

In some cases FORCE is unable to successfully prune past a certain sparsity. In particular, at some point in the pruning schedule, FORCE begins to assign all parameters a saliency score of 0, thus providing no basis for pruning, with the consequence being that the algorithm simply returns a dense network. This happens, for example, at the most extreme sparsities in VGG19, ResNet18, and MobileNetV2. In these cases, the test accuracy for FORCE is reported as equal to that of random guessing, since the algorithm cannot provide a trainable network at the given sparsity, but this is



denoted with a dashed lined to indicate that no network of the specified sparsity was actually tested.

We conjecture that this phenomenon is a result of *throughput collapse*<sup>9</sup> - that is, once the algorithm fully prunes *all branches* of communication at some point in the network, though we did not investigate this further. We note that in investigating these collapses, we also tried doubling the number of pruning steps from 60 to 120, but this did not avoid the problem.

## D.2. Fixup ResNet

Neither FORCE nor SynFlow can be applied without modification to FixupResNet110. As above, this failure is due to the fact that both algorithms assign a saliency score of 0 to all parameters and consequently have no basis on which to prune.

In particular, at initialization, the only non-zero gradients of both the training loss  $\mathcal{L}$  and SynFlow’s objective function  $\mathcal{R}$ , are in the network’s final layer, where the weights themselves are initialized as 0. The saliency scores in both FORCE and SynFlow are obtained via the elementwise multiplication of the parameter matrices with their gradients, and thus are 0 in all cases.

## E. Experimental Details

### E.1. Code and Implementation

We implemented Force<sup>10</sup> and SynFlow<sup>11</sup> using the code published by the respective authors, adapted to include any additional architectures used in our experiments.

The code used to run experiments with DCTpS networks is included as additional supplementary material.

### E.2. Model Architectures

Standard implementations of network architectures used here are taken from the following sources:

- ResNet50, ResNet18 and VGG19, as implemented in (de Jorge et al., 2021). See [the FORCE Github Repo](#)
- MobileNetV2 from the authors’ published code [here](#).
- FixupResNet110 from the authors’ published code [here](#).

### E.3. Parameter Breakdown by Architecture

See Table 2 for a breakdown of the prunable/non-prunable parameter totals in each of the architectures used in our experiments.

<sup>9</sup>This is equivalent to layer collapse when there is only a single feedforward connection at each layer.

<sup>10</sup><https://github.com/naver/force>

<sup>11</sup><https://github.com/ganguli-lab/Synaptic-Flow>

Table 2. Division of total parameters between weights (pruned) and bias and/or batchnorm (BN) parameters (not pruned) in the architectures used in our experiments, with 10 output classes.

	Weights	Bias & BN	Total
ResNet50	23467712	53130	23520842
VGG19	20024000	11018	20035018
MobileNetV2	2261824	35098	2296922
FixupResNet110	1719856	282	1720138
ResNet18	11164352	9610	11173962

### E.4. Pruning hyperparameters

See Table 3 for the hyperparameters used when applying FORCE and SynFlow.

Table 3. Pruning hyperparameters for FORCE and SynFlow. C10, C100, and TI stand for CIFAR10, CIFAR100, and Tiny Imagenet respectively.

	FORCE	SynFlow
Prune Steps	60	100
# Batches	1 (C10), 10 (C100), 20 (TI)	N/A
Schedule	exp	exp

Table 4. Training hyperparameters used for experiments in Section 5 and Appendix A. Note that for training DCTpS networks with SGD, a base learning rate of 0.03 was used instead of 0.1. For experiments with Lenet-5 (only performed with Adam on CIFAR10) batch size was 64 and total epochs was 160.

	Adam	SGD
Epochs	200	200
Batch Size	128	128
Learning Rate (LR)	0.001	0.1
Momentum	N/A	0.9
LR Decay Epochs	N/A	120, 160
LR Drop factor	N/A	0.1
Weight Decay	$5 \times 10^{-4}$	$5 \times 10^{-4}$

---

### E.5. Training Details

See Table 4 for the training hyperparameters used in our experiments in Section 5 and Appendix A. On CIFAR10 and CIFAR100 (Krizhevsky et al., 2009), 10% of the training data is withheld as a validation set. The model with the maximum validation accuracy is selected as our final model, to be evaluated on the test set. In the case of Tiny Imagenet (Wu et al., 2017), where there are no labels for the test set, the maximum validation accuracy obtained during training is reported. All experiments were run using Adam, except for those in Appendix A.3 in which SGD with momentum was used.