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

Ilan Price^{1,2} Jared Tanner^{1,2}

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.

¹Mathematical Institute, University of Oxford, Oxford, UK

²The Alan Turing Institute, London, UK. 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 PaI 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.

An alternative approach to training sparse networks from scratch is Dynamic Sparse Training (DST) (Mocanu et al., 2018; Mostafa & Wang, 2019). In DST methods, the network is initialized as sparse according to some heuristic

Table 1. Summary of state-of-the-art PaI and DST 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 average 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). If this cost differs between the forward and backward pass, the larger of the two is included. The extra factor of $\frac{1}{\Delta T}mn$ for RigL comes from the necessity to compute full gradients every ΔT steps. We report the drop in accuracy for the case of ResNet50 applied to CIFAR100 at different global sparsities, relative to the dense baseline. Quantities which decrease from $\mathcal{O}(C)$ to $\mathcal{O}(D)$ during training are denoted by $\mathcal{O}(C) \xrightarrow{t} \mathcal{O}(D)$.

| $P =$ | Δ Accuracy | | | Computational Cost | | | Network Size on Device | | |
|--------------|-------------------|--------|--------|--------------------|--|-------------------------------|------------------------|--|-------------------|
| | 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 | +0.8% | -7.1% | -64.8% | 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% | $\mathcal{O}(mnk)$ | $\mathcal{O}(pmn)$ | $\mathcal{O}(pmn)$ | $\mathcal{O}(N)$ | $\mathcal{O}(PN)$ | $\mathcal{O}(PN)$ |
| RigL (ERK) | +0.4% | -16.8% | -65.7% | 0 | $\mathcal{O}(pmn + \frac{1}{\Delta T}mn)$ | $\mathcal{O}(pmn)$ | $\mathcal{O}(PN)$ | $\mathcal{O}(PN)$ | $\mathcal{O}(PN)$ |
| DCTpS | -5.8% | -15% | -22.8% | 0 | $\mathcal{O}(q \log q + pmn)$ | $\mathcal{O}(q \log q + pmn)$ | $\mathcal{O}(PN)$ | $\mathcal{O}(PN)$ | $\mathcal{O}(PN)$ |

(thereby also avoiding any initialization and storage of the full network), but the topology - the support set of the non-zero parameters - of the network is updated during training, along with the values of the weights themselves.

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

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 at which accuracy rapidly collapses to random guessing. The primary improvement 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. Our approach and results can be summarised 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 in 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 density 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.
- Combining DCTpS layers with state-of-the-art Dynamic Sparse Training (DST) techniques yields even further improvements in accuracy at all densities - often in the region of 3%-7%.

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: $G(\mathbf{w}) = \left| \frac{\partial \mathcal{L}(\mathbf{w})}{\partial \mathbf{w}} \odot \mathbf{w} \right|$, 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 $G(\mathbf{w}) = \left| \frac{\partial \mathcal{L}(\bar{\mathbf{w}})}{\partial \mathbf{w}} \odot \mathbf{w} \right|$, 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^{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.

PaI methods are also sometimes referred to as static sparse training (SST), since the sparse support remains fixed after initialization. DST methods (Mocanu et al., 2018; Mostafa & Wang, 2019), on the other hand, modify the support set during training, subject to maintaining a fixed sparsity level overall. DST methods are characterised by the criteria according to which they prune some weights and ‘re-grow’/activate others during training, and also whether the layer-wise sparsity levels remain fixed during training, or if connections pruned in one layer can be replaced by additional connections in other layers. The current state-of-the-art for DST is RigL (Evcı et al., 2020), which maintains a fixed layer-wise sparsity distribution, prunes parameters with the smallest magnitude, and re-activates weights with the largest-magnitude gradients.

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 SST methods is due to their layer-wise distribution of the number of trainable parameters, rather than the particular locations of the trainable parameters within a layer. This somewhat calls into question the role of the proposed saliency metrics used to score the importance of each parameter individually. In contrast, while RigL is also shown to be influenced by the layer-wise sparsity distribution (Evcı et al., 2020), the premise of DST is precisely that the locations of the trainable parameters within each layer matter fundamentally. Either way, further understanding of, and heuristics for, the ideal layer-wise parameter allocations would be complementary and directly beneficial to the aforementioned PaI and DST methods, as well as ‘DCT plus Sparse’ presented in Section 4.

3. Restricting Network Weights to Random Subspaces

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

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

where 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 (w being k -sparse, with support set \mathcal{S}), such as those generated by PaI methods, represents the specific case when $d = 0$, 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 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 θ_0 .

The model (2) was explored in (Li et al., 2018) where they showed it is possible to *randomly* draw the offset 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 θ . In their work, 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) 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 random low-dimensional subspace 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 \mathcal{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.

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.

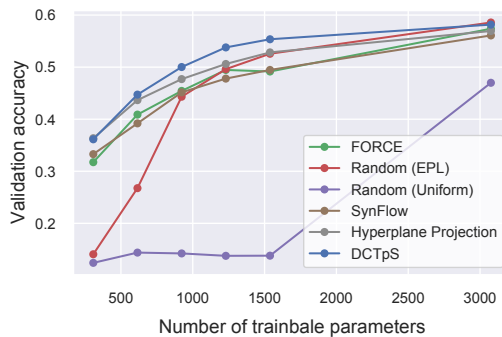


Figure 1. Different subspace selection methods applied to Lenet-5 (LeCun et al., 1998) and trained on CIFAR10. We report maximum validation accuracy, averaged over three runs, at each subspace dimension (number of trainable parameters).

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 parameters¹. 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 θ). 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 parameter² α , similar to batch normalization (Ioffe & Szegedy, 2015), we consider $W = \alpha D + S$.

In order to maintain the low network size on device and

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

²We note that the inclusion of an α 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 with almost no overhead. See the Supp. Mat. for experiments without an α parameter.

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) matrix³ 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 SynFlow, 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 select their locations within each weight tensor 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.

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

5. Experiments

SynFlow (Tanaka et al., 2020) and FORCE (de Jorge et al., 2021) have 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 FORCE⁴, as well as RigL (ERK), the current state-of-the-art DST method. A full description of all experimental setups and hyperparameters is included in the Supplementary Material, and our code is available at github.com/IlanPrice/DCTpS. 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 indicates where an algorithm breaks down and is thus unable to prune the network to the specified sparsity.

Combining DCTpS with RigL: Though DCTpS networks have been presented in the framework of static sparse training, they can straightforwardly be combined with DST methods like RigL, which are then used to iteratively update the support set of the DCTpS layers’ sparse S matrices. We test this in the experiments below, with positive results.

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 SST 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 Random (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, Random (EPL), which uses the 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 Random (EPL) heuristic significantly outperforms Random (uniform) and other heuristics (see Supp. Mat.), and even matches state-of-the-art pruning methods for densities greater than $\sim 1\%$, see e.g. Figure 2.

5.1. Lenet-5

We first consider the small Lenet-5 architecture on CIFAR10 so as to compare sparse network methods against the ‘Hyperplane Projection’ method of (Li et al., 2018), which is

⁴A revised version of (de Jorge et al., 2021) includes a closely related variant of FORCE, called Iter-SNIP with very similar results. It suffices to compare our method to FORCE.

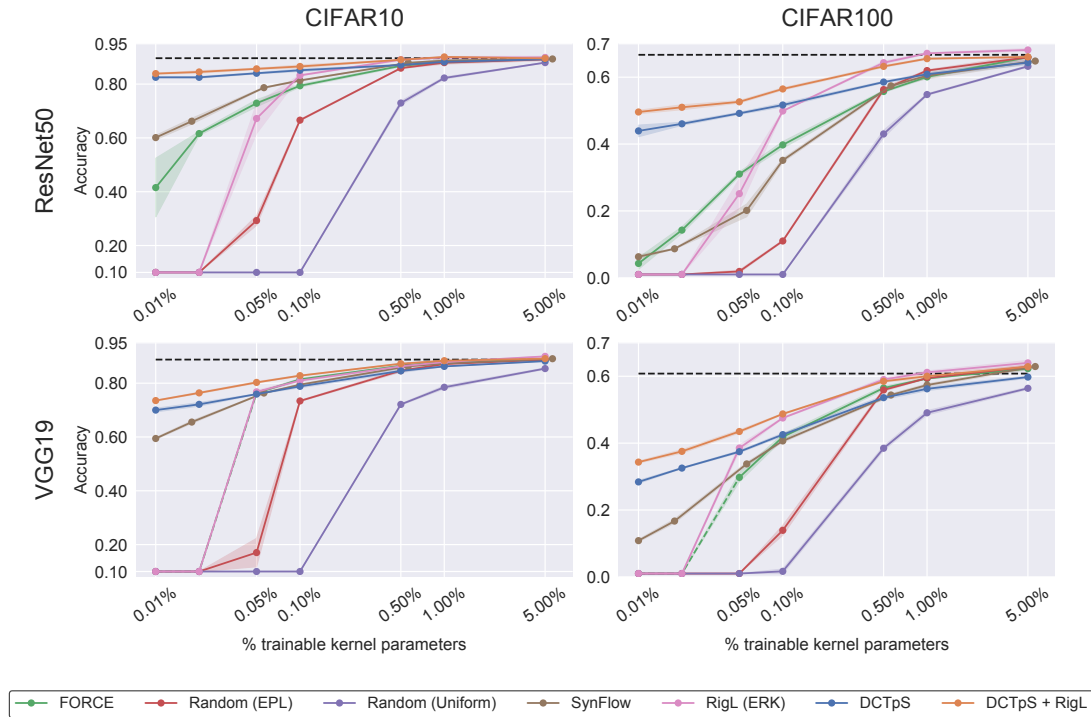


Figure 2. Test accuracy on CIFAR10 and CIFAR100 datasets using sparse ResNet50 and VGG19 architectures. DCT plus Sparse (DCTpS) networks (with EPL parameter allocation) are compared with FORCE, SynFlow, RigL, and random pruning methods.

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 shows that the offset can be deterministic and the hyperplane sparse and axis-aligned as in (3).

5.2. ResNet50 and VGG19 on 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 test 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. Random (uniform) and Random (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 PaI methods.

Below 0.5%, DCTpS has superior or equal test accuracy compared to both SynFlow⁵ 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%. With the introduction of dynamic sparsity, RigL achieves superior accuracy to all static sparse approaches at higher densities (ranging from 5% to 0.5%), after which it is overtaken by DCTpS. Applying RigL to the sparse matrices in DCTpS networks yields significant improvements over static sparse DCTpS networks; for example, considering ResNet50 applied to CIFAR100, DCTpS with RigL achieves $\approx 50\%$ accuracy at 0.01% density, as compared with 44% with static DCTpS. DCTpS combined with RigL achieves accuracy comparable to or better than all methods, including RigL, at densities of 1% or below, again with the most significant margins at the most extreme sparsities.

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

⁵We 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.

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. Once again RigL outperforms static sparse approaches at densities ranging from 5% to 0.1%, but collapses below this level. Combining RigL with DCTpS is comparable to or better than all other methods at all densities, and significantly outperforms static DCTpS (by 3% - 6% at all densities on CIFAR100, for example).

One important point to underscore is that in these experiments, as is typical in prior PaI experiments (de Jorge et al., 2021), the reported sparsity refers only to the 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 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 and RigL 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 those 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. RigL’s accuracy drops significantly below 1% density, but combining DCTpS with RigL achieves the best accuracy of all methods at all densities.

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’, practi-

cally 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 α parameter to 0 or to the appropriate scaling factor. We use the code provided by the authors⁶ to obtain the baseline and random pruning results, and create DCTpS Fixup ResNets by simply replacing the Linear and Convolutional layers with their corresponding DCTpS variants, and initializing appropriately.

Figure 3 (right) illustrates the test accuracy of DCTpS, RigL, Random (EPL) and Random (uniform). 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 Random (EPL) by more than 20%. DCTpS outperforms RigL below 0.5% density, and DCTpS with RigL matches or outperforms all other methods at all densities.

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 p 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 $p < \mathcal{O}(\log q/r)$ where $r = \min(m, n)$, 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

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

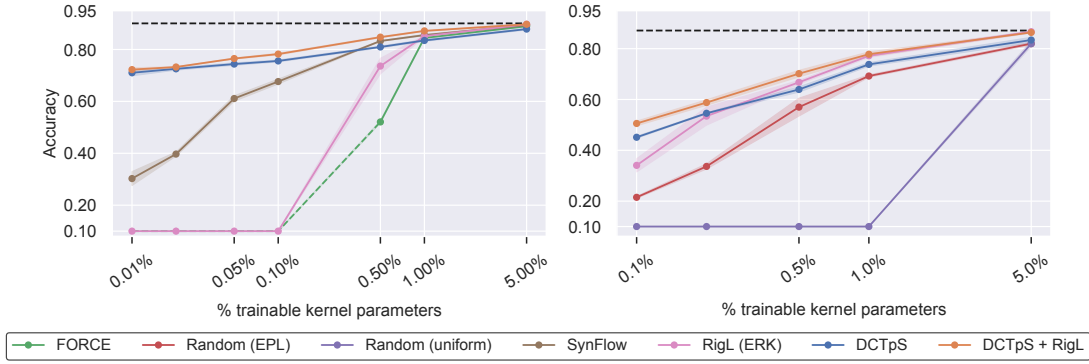


Figure 3. Left: Comparison of DCTpS against FORCE, SynFlow and RigL on CIFAR10 with MobileNetV2 (Sandler et al., 2018). Right: Comparison of DCTpS with RigL and Random Pruning methods on CIFAR10 with FixupResNet110 (Zhang et al., 2019).

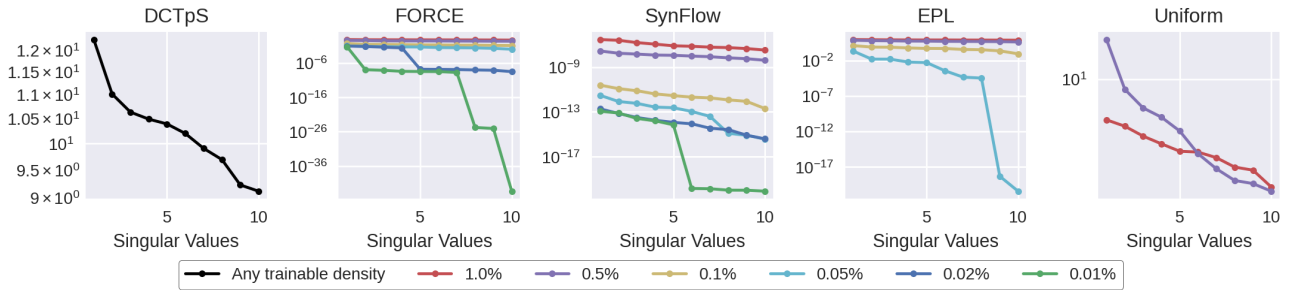


Figure 4. Spectrum of the Jacobian of Resnet50 on CIFAR10, pruned to different sparsities with varies methods, at initialization. If a curve does not appear in a given plot, it means that the spectrum was identically zero for that density. The DCTpS plot shows only one curve since its Jacobian does not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

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 Random (uniform) collapses to random guessing as the density decreases from 0.5% to 0.1%, and for Random (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 $> 10^{10}$ 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 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



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.

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 Random (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 distribution of nonzeros generated by PaI methods was recently investigated in depth by (Frankle et al., 2021); and

in particular, the value of carefully selecting the locations of trainable parameters with PaI methods. It was observed that given a particular sparsity pattern identified by PaI methods, the location of a layer’s nonzeros 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 number of trainable parameters per layer rather than which entries within the layers 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 Random (EPL).

7. Conclusions and Further Work

In this work we have shown that adding a layer-wise offset to sparse subspace training significantly improves test 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. These results are further improved when DST methods like RigL are applied to DCTpS networks.

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 α 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.
- 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>.
- Evci, U., Gale, T., Menick, J., Castro, P. S., and Elsen, E. Rigging the lottery: Making all tickets winners. In *International Conference on Machine Learning*, pp. 2943–2952. PMLR, 2020.
- 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=Iq-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.
- 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.
- 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=B1VZqjAcYX>.
- Li, C., Farkhoor, 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.
- 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):1–12, 2018.
- Mostafa, H. and Wang, X. Parameter efficient training of deep convolutional neural networks by dynamic sparse reparameterization. In *International Conference on Machine Learning*, pp. 4646–4655. PMLR, 2019.
- 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>.
- 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>.