

i-SpaSP: Structured Neural Pruning via Sparse Signal Recovery

Cameron R. Wolfe

Department of Computer Science, Rice University, Houston, TX, USA.

CRW13@RICE.EDU

Anastasios Kyrillidis

Department of Computer Science, Rice University, Houston, TX, USA.

ANASTASIOS@RICE.EDU

Editors: R. Firoozi, N. Mehr, E. Yel, R. Antonova, J. Bohg, M. Schwager, M. Kochenderfer

Abstract

We propose a novel, structured pruning algorithm for neural networks—the **iterative, Sparse Structured Pruning** algorithm, dubbed as i-SpaSP. Inspired by ideas from sparse signal recovery, i-SpaSP operates by iteratively identifying a larger set of important parameter groups (e.g., filters or neurons) within a network that contribute most to the residual between pruned and dense network output, then thresholding these groups based on a smaller, pre-defined pruning ratio. For both two-layer and multi-layer network architectures with ReLU activations, we show the error induced by pruning with i-SpaSP decays polynomially, where the degree of this polynomial becomes arbitrarily large based on the sparsity of the dense network’s hidden representations. In our experiments, i-SpaSP is evaluated across a variety of datasets (i.e., MNIST, ImageNet, and XNLI) and architectures (i.e., feed forward networks, ResNet34, MobileNetV2, and BERT), where it is shown to discover high-performing sub-networks and improve upon the pruning efficiency of provable baseline methodologies by several orders of magnitude. Put simply, i-SpaSP is easy to implement with automatic differentiation, achieves strong empirical results, comes with theoretical convergence guarantees, and is efficient, thus distinguishing itself as one of the few computationally efficient, practical, and provable pruning algorithms.

Keywords: Neural Network Pruning, Greedy Selection, Sparse Signal Recovery

Source Code: <https://github.com/wolfecameron/i-SpaSP>

Appendix: <https://arxiv.org/abs/2112.04905>

1. Introduction

Background. Neural network pruning has garnered significant recent interest (Frankle and Carbin, 2018; Liu et al., 2018; Li et al., 2016), as obtaining high-performing sub-networks from larger, dense networks enables a reduction in the computational and memory overhead of neural network applications (Han et al., 2015a,b). Many popular pruning techniques are based upon empirical heuristics that work well in practice (Frankle et al., 2019; Luo et al., 2017; He et al., 2017). Generally, these methodologies introduce some notion of “importance” for network parameters (or groups of parameters) and eliminate parameters with negligible importance.

The empirical success of pruning methodologies inspired the development of pruning algorithms with theoretical guarantees (Baykal et al., 2019; Liebenwein et al., 2019; Mussay et al., 2019; Baykal et al., 2018; Ramanujan et al., 2019). Among such work, greedy forward selection (GFS) (Ye et al., 2020; Ye et al., 2020)—inspired by the Frank-Wolfe algorithm (Frank et al., 1956)—differentiated itself as a methodology that performs well in practice and provides theoretical guarantees. However, GFS is inefficient in comparison to popular pruning heuristics.

This Work. We leverage greedy selection (Needell and Tropp, 2009; Khanna and Kyriallidis, 2018) to develop a structured pruning algorithm that is provable, practical, and efficient. This algorithm, called **iterative, Sparse Structured Pruning (i-SpaSP)**, iteratively estimates the most important parameter groups¹ within each network layer, then thresholds this set of parameter groups based on a pre-defined pruning ratio—a similar procedure to the CoSAMP algorithm for sparse signal recovery (Needell and Tropp, 2009). Theoretically, we show for two and multi-layer networks that *i*) the output residual between pruned and dense networks decays polynomially with respect to the size of the pruned network and *ii*) the order of this polynomial increases as the dense network’s hidden representations become more sparse.² In experiments, we show that i-SpaSP is capable of discovering high-performing sub-networks across numerous different models (i.e., two-layer networks, ResNet34, MobileNetV2, and BERT) and datasets (i.e., MNIST, ImageNet, and XNLI). i-SpaSP is simple to implement and significantly improves upon the runtime of GFS variants.

2. Preliminaries

Notation. Vectors and scalars are denoted with lower-case letters. Matrices and certain constants are denoted with upper-case letters. Sets are denoted with upper-case, calligraphic letters (e.g., \mathcal{G}) with set complements \mathcal{G}^c . We denote $[n] = \{0, 1, \dots, n\}$. For $x \in \mathbb{R}^N$, $\|x\|_p$ is the ℓ_p vector norm. x_s is the s largest-valued components of x , where $|x| \geq s$. $\text{supp}(x)$ returns the support of x . For index set \mathcal{G} , $x|_{\mathcal{G}}$ is the vector with non-zeros at the indices in \mathcal{G} . For $X \in \mathbb{R}^{m \times n}$, $\|X\|_F$ and X^\top represent the Frobenius norm and transpose of X . The i -th row and j -th column of X are given by $X_{i,:}$ and $X_{:,j}$, respectively. $\text{rsupp}(X)$ returns the row support of X (i.e., indices of non-zero rows). For index set \mathcal{G} , $X_{\mathcal{G},:}$ and $X_{:, \mathcal{G}}$ represent row and column sub-matrices, respectively, that contain rows or columns with indices in \mathcal{G} . $\mu(X) : \mathbb{R}^{m \times n} \rightarrow \mathbb{R}^m$ sums over columns of a matrix (i.e., $\mu(X) = \sum_i X_{:,i}$), while $\text{vec}(X) : \mathbb{R}^{m \times n} \rightarrow \mathbb{R}^{mn}$ stacks columns of a matrix. $X \in \mathbb{R}^{m \times n}$ is p -row-compressible with magnitude $R \in \mathbb{R}^{\geq 0}$ if $|\mu(X)|_{(i)} \leq \frac{R}{i^p}$, $\forall i \in [m]$, where $|\cdot|_{(i)}$ denotes the i -th sorted vector component (in magnitude). Lower p values indicate a nearly row-sparse matrix and vice versa.

Network Architecture. Our analysis primarily considers two-layer, feed forward networks³:

$$f(X, W) = W^{(1)} \cdot \sigma(W^{(0)} \cdot X) \quad (1)$$

The network’s input, hidden, and output dimensions are given by d_{in} , d_{hid} , and d_{out} . $X \in \mathbb{R}^{d_{in} \times B}$ stores the full input dataset with B examples. $W^{(0)} \in \mathbb{R}^{d_{hid} \times d_{in}}$ and $W^{(1)} \in \mathbb{R}^{d_{out} \times d_{hid}}$ denote the network’s weight matrices. $\sigma(\cdot)$ denotes the ReLU activation function and $H = \sigma(W^{(0)} \cdot X)$ stores the network hidden representations across the dataset. We also extend our analysis to multi-layer networks with similar structure; see Appendix ?? for more details.

3. Related Work

Pruning. Neural network pruning strategies can be roughly separated into structured (Han et al., 2016; Li et al., 2016; Liu et al., 2017; Ye et al., 2020; Ye et al., 2020) and unstructured (Evci et al.,

1. “Parameter groups” refers to the minimum structure used for pruning (e.g., neurons or filters).

2. To the best of the authors’ knowledge, we are the first to provide theoretical analysis showing that the quality of pruning depends upon the sparsity of representations within the dense network.

3. Though (1) has no bias, a bias term could be implicitly added as an extra element within the input and weight matrices.

2019, 2020; Frankle and Carbin, 2018; Han et al., 2015) variants. Structured pruning, as considered in this work, prunes parameter groups instead of individual weights, allowing speedups to be achieved without sparse computation (Li et al., 2016). Empirical heuristics for structured pruning include removing parameter groups with low ℓ_1 norm (Li et al., 2016; Liu et al., 2017), measuring the gradient-based sensitivity of parameter groups (Baykal et al., 2019; Wang et al., 2020; Zhuang et al., 2018), preserving network output (He et al., 2017; Luo et al., 2017; Yu et al., 2017), and more (Suau et al., 2020; Chin et al., 2019; Huang and Wang, 2018; Molchanov et al., 2016). Pruning typically follows a three-step process of pre-training, pruning, and fine-tuning (Li et al., 2016; Liu et al., 2018), where pre-training is typically the most expensive component (Chen et al., 2020; You et al., 2019).

Provable Pruning. Empirical pruning research inspired the development of theoretical foundations for network pruning, including sensitivity-based analysis (Baykal et al., 2019; Liebenwein et al., 2019), coreset methodologies (Mussay et al., 2019; Baykal et al., 2018), random network pruning analysis (Malach et al., 2020; Orseau et al., 2020; Pensia et al., 2020; Ramanujan et al., 2019), and generalization analysis (Zhang et al., 2021). GFS—analyzed for two-layer (Ye et al., 2020; Wolfe et al., 2021) and multi-layer (Ye et al., 2020) networks—was one of the first pruning methodologies to provide both strong empirical performance and theoretical guarantees.

Greedy Selection. Greedy selection efficiently discovers approximate solutions to combinatorial optimization problems (Frank et al., 1956). Many algorithms and frameworks for greedy selection exist; e.g., Frank-Wolfe (Frank et al., 1956), sub-modular optimization (Nemhauser et al., 1978), CoSAMP (Needell and Tropp, 2009), and iterative hard thresholding (Khanna and Kyrillidis, 2018). Frank-Wolfe has been used within GFS and to train deep neural networks (Bach, 2014; Pokutta et al., 2020), thus forming a connection between greedy selection and deep learning. Furthering this connection, we leverage CoSAMP (Needell and Tropp, 2009) to formulate our proposed methodology.

4. Methodology

i-SpaSP is formulated for two-layer networks in Algorithm 1, where the pruned model size s and total iterations T are fixed. \mathcal{S} stores active neurons in the pruned model, which is refined over iterations.

Why does this work? Each iteration of i-SpaSP follows a three-step procedure in Algorithm 1:

- Step I:** Compute neuron “importance” Y given the current residual matrix V .
- Step II:** Identify s neurons within the combined set of important and active neurons (i.e., $\Omega \cup \mathcal{S}$) with the largest-valued hidden representations.
- Step III:** Update V with respect to the new pruned model estimate.

We now provide intuition regarding the purpose of each individual step within i-SpaSP.

Estimating Importance. Y_{ij} is the importance of hidden neuron i with respect to dataset example j . We can characterize the discrepancy between pruned and dense network output U' and U as:

$$\mathcal{L}(U, U') = \frac{1}{2} \|W^{(1)} \cdot H - U'\|_F^2. \quad (2)$$

Considering U' fixed, $\nabla_H \mathcal{L}(U, U') = (W^{(1)})^\top \cdot V$. As such, if Y_{ij} is a large, positive (negative) value, decreasing (increasing) H_{ij} will decrease the value of \mathcal{L} locally. Then, because H is non-negative and cannot be modified via pruning, one can realize that the best methodology of minimizing (2) is including neurons with large, positive importance values within \mathcal{S} , as in Algorithm 1.

Algorithm 1 i-SpaSP for Two-Layer Networks**Parameters:** $T, s; \mathcal{S} := \emptyset; t := 0$

compute hidden representation

$$H = \sigma(W^{(0)} \cdot X)$$

$$h = \mu(H)$$

compute dense network output

$$U = W^{(1)} \cdot H$$

$$V = U$$

while $t < T$ **do**

$$t = t + 1$$

Step I: Estimating Importance

$$Y = (W^{(1)})^\top \cdot V$$

$$y = \mu(Y)$$

$$\Omega = \text{supp}(y_{2s})$$

Step II: Merging and Pruning

$$\Omega^* = \Omega \cup \mathcal{S}$$

$$b = h|_{\Omega^*}$$

$$\mathcal{S} = \text{supp}(b_s)$$

Step III: Computing New Residual

$$V = U - W_{:, \mathcal{S}}^{(1)} \cdot H_{\mathcal{S}, :}$$

end# return pruned model with neurons in \mathcal{S}

$$\text{return } \{W_{\mathcal{S}, :}^{(0)}, W_{:, \mathcal{S}}^{(1)}\}$$

Algorithm 2 for a PyTorch-style example. Automatic differentiation simplifies importance estimation and allows it to be run on a GPU, making the implementation efficient and parallelizable. Because the remainder of the pruning process only leverages basic sorting and set operations, the overall implementation of i-SpaSP is both simple and efficient.

Other Architectures. Algorithm 2 can be easily generalized to more complex network modules (i.e., beyond feed-forward layers) using automatic differentiation. Notably, convolutional filters or attention heads can be pruned using the importance estimation from Algorithm 2 if $\text{sum}(\cdot)$ is performed over both batch and spatial dimensions. Furthermore, i-SpaSP can be used to prune multi-layer networks by greedily pruning each layer of the network from beginning to end.

Large-Scale Datasets. Algorithm 1 assumes the entire dataset is stored within X . For large-scale experiments, such an approach is not tractable. As such, we redefine X within experiments to contain a subset or mini-batch of data from the full dataset, allowing the pruning process to be performing in an approximate—but computationally tractable—manner. To improve the quality of this approximation, a new mini-batch is sampled during each i-SpaSP iteration, but the size of such mini-batches becomes a hyperparameter of the pruning process.⁴

Merging and Pruning. The $2s$ most-important neurons—based on $\mu(Y)$ components—are selected and merged with \mathcal{S} , allowing a larger set of neurons (i.e., more than s) to be explored. From here, s neurons with the largest-valued components in $\mu(H)$ are sub-selected from this combined set to form the next pruned model estimate. Because hidden representation values are not considered in importance estimation, performing this two-step merging and pruning process ensures neurons within \mathcal{S} have both large hidden activation and importance values, which together indicate a meaningful impact on network output.

Computing the New Residual. The next pruned model estimate is used to re-compute V , which can be intuitively seen as updating U' in (2). As such, importance in Algorithm 1 is based on the current pruned and dense network residual and (2) is minimized over successive iterations.

4.1. Implementation

Automatic Differentiation. Because $Y = \nabla_H \mathcal{L}(U, U')$, importance within i-SpaSP can be computed efficiently using automatic differentiation (Paszke et al., 2017; Abadi et al., 2015); see

4. Both GFS (Ye et al., 2020) and multi-layer GFS (Ye et al., 2020) adopt a similar mini-batch approach during pruning.

```

H.requires_grad := True
with torch.no_grad():
    prune_out := prune_layer(Hs,:)
dense_out := dense_layer(H)
obj := sum(1/2(dense_out - prune_out)^2)
obj.backward()
importance := sum(H.grad, dim = 0)
return importance
    
```

Algorithm 2: i-SpaSP importance computation via automatic differentiation.

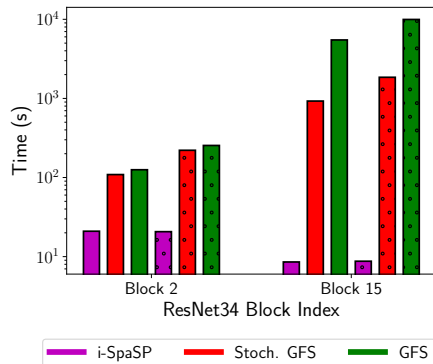


Figure 1: Runtime of pruning ResNet34 blocks with i-SpaSP and GFS variants to 20% (i.e., plain) or 40% (i.e., dotted) of filters.

4.2. Computational Complexity and Runtime Comparisons

Denote the complexity of matrix-matrix multiplication as ξ . The complexity of pruning a network layer with T iterations of i-SpaSP is $\mathcal{O}(T\xi + Td_{hid} \log(d_{hid}))$. GFS has a complexity of $\mathcal{O}(s\xi d_{hid})$, as it adds a single neuron to the (initially empty) network layer at each iteration by exhaustively searching for the neuron that minimizes training loss. Though later GFS variants achieve complexity of $\mathcal{O}(s\xi)$ (Ye et al., 2020), i-SpaSP is more efficient in practice because the forward pass (i.e. $\mathcal{O}(\xi)$) dominates the pruning procedure and $T \ll s$ (e.g., $T = 20$ in Section 6).

As a practical runtime comparison, we adopt a ResNet34 model (He et al., 2015) and measure wall-clock pruning time⁵ with i-SpaSP and GFS. We use the public implementation of GFS (Ye, 2021) and test both stochastic and vanilla variants⁶. i-SpaSP uses settings from Section 6, and selected ResNet blocks are pruned to ratios of 20% or 40% of original filters; see Figure 1. i-SpaSP significantly improves upon the runtime of GFS variants; e.g., i-SpaSP prunes Block 15 in roughly 10 seconds, while GFS takes over 1000 seconds in the best case. Furthermore, unlike GFS, the runtime of i-SpaSP is not sensitive to the size of the pruned network (i.e., wall-clock time is similar for ratios of 20% and 40%), though i-SpaSP does prune later network layers faster than earlier layers.

5. Theoretical Results

Proofs are deferred to Appendix ???. The dense network is pruned from d_{hid} to s neurons via i-SpaSP. We assume that $W^{(1)}$ satisfies the restricted isometry property (RIP) (Candes and Tao, 2006):

Assumption 1 (Restricted Isometry Property (RIP) (Candes and Tao, 2006)) Denote the r -th restricted isometry constant as δ_r and assume $\delta_r \leq 0.1$ for $r = 4s$.⁷ Then, $W^{(1)}$ satisfies the RIP with constant δ_r if $(1 - \delta_r)\|x\|_2^2 \leq \|W^{(1)} \cdot x\|_2^2 \leq (1 + \delta_r)\|x\|_2^2$ for all $\|x\|_0 \leq r$.

5. We prune the 2nd (64 channels) and 15th (512 channels) convolutional blocks. We choose blocks in different network regions to view the impact of channel and spatial dimension on pruning efficiency.
 6. Vanilla GFS exhaustively searches neurons within each iteration, while the stochastic variant randomly selects 50 neurons to search per iteration.
 7. This is a numerical assumption adopted from (Needell and Tropp, 2009), which holds for Gaussian matrices of size $\mathbb{R}^{m \times n}$ when $m \geq \mathcal{O}(r \log(\frac{n}{r}))$.

No assumption is made upon $W^{(0)}$. We hypothesize that Assumption 1 is mild due to properties like semi or quarter-circle laws that bound the eigenvalues symmetric, random matrices within some range (Edelman and Wang, 2013), but we leave the formal verification of this assumption as future work. We define $H = \sigma(W^{(0)} \cdot X) \in \mathbb{R}^{d_{hid} \times B}$, which can be theoretically reformulated as $H = Z + E$ for s -row-sparse Z and arbitrary E . From here, we can show the following about the residual between pruned and dense network hidden representations after t iterations of Algorithm 1.

Lemma 1 *If Assumption 1 holds, the pruned approximation to H after t iterations of Algorithm 1, $H_{S_t, \cdot}$ is s -row-sparse and satisfies the following inequality:*

$$\|\mu(H - H_{S_t, \cdot})\|_2 \leq (0.444)^t \|\mu(H)\|_2 + \left(14 + \frac{7}{\sqrt{s}}\right) \|\mu(E)\|_1$$

Going further, Lemma 1 can be used to bound the residual between pruned and dense network output.

Theorem 2 *Let $U = W^{(1)} \cdot H$ and $U' = W_{:, S_t}^{(1)} \cdot H_{S_t, \cdot}$ denote pruned and dense network output, respectively. $V_t = U - U'$ stores the residual between pruned and dense network output over the entire dataset. If Assumption 1 holds, we have the following at iteration t of Algorithm 1:*

$$\|V_t\|_F \leq \|W^{(1)}\|_F \cdot \left((0.444)^t \|\mu(H)\|_2 + \left(14 + \frac{7}{\sqrt{s}}\right) \|\mu(E)\|_1 \right)$$

Because $\|\mu(H)\|_2$ decays linearly in Theorem 2, $\|\mu(E)\|_1$ dominates the above expression for large t . By assuming H is row-compressible, we can derive a bound on $\|\mu(E)\|_1$.

Lemma 3 *Assume H is p -row-compressible with magnitude R , where $H = Z + E$ for s -row-sparse Z and arbitrary E . Then, $\|\mu(E)\|_1 \leq R \cdot \frac{s^{-\frac{1}{p}}}{\frac{1}{p}-1}$.*

Lemma 3 can then be combined with Theorem 2 to bound the error due to pruning via i-SpaSP.

Theorem 4 *Assume Algorithm 1 is run for a sufficiently large number of iterations t . If Assumption 1 holds and H is p -row-compressible with factor R , the output residual between the dense network and the pruned network discovered via i-SpaSP can be bounded as follows:*

$$\|V_t\|_F \leq \mathcal{O} \left(\frac{s^{\frac{1}{2}-\frac{1}{p}} p (2\sqrt{s} + 1)}{1 - p} \right).$$

Proof *This follows directly from substituting Lemma 3 into Theorem 2, assuming t is large enough such that $(0.444)^t \|W^{(1)}\|_F \|\mu(H)\|_2 \approx 0$, and factoring out constants in the resulting expression. ■*

Theorem 4 indicates that the quality of the pruned network is dependent upon s and p . Intuitively, one would expect that lower values of p (implying sparser H) would make pruning easier, as neurons corresponding to zero rows in H could be eliminated without consequence. This trend is observed exactly within Theorem 4; e.g., for $p = \{\frac{3}{4}, \frac{1}{2}, \frac{1}{4}\}$ we have $\|V_t\|_F \leq \{\mathcal{O}(s^{-\frac{1}{3}}), \mathcal{O}(s^{-1}), \mathcal{O}(s^{-3})\}$, respectively. *To the best of the authors' knowledge, our work is the first to theoretically characterize pruning error with respect to sparsity properties of network hidden representations.* This bound can also be extended to similarly-structured (see Appendix ??), multi-layer networks:

Theorem 5 Consider an L -hidden-layer network with weight matrices $\{W^{(0)}, \dots, W^{(L)}\}$ and hidden representations $\{H^{(1)}, \dots, H^{(L)}\}$. The hidden representations of each layer are assumed to have dimension d for simplicity. We define $H^{(\ell)} = \sigma(W^{(\ell)} \cdot H^{(\ell-1)})$, where $H^{(1)} = \sigma(W^{(0)} \cdot X)$ and $H^{(L)} = W^{(L)} \cdot H^{(L-1)}$. We assume all weight matrices other than $W^{(0)}$ obey Assumption 1 and all hidden representations other than $H^{(L)}$ are p -row-compressible. i -SpaSP is applied greedily to prune each network layer, in layer order, from d to s hidden neurons. Given sufficient iterations t , the residual between pruned and dense multi-layer network output behaves as:

$$\|V_t^{(L)}\|_F \leq \mathcal{O} \left(\sum_{i=1}^L \left(14 + \frac{7}{\sqrt{s}} \right)^{L-i+1} \left(\|W^{(L)}\|_F \prod_{j=1}^{L-i} \|\text{vec}(W^{(j)})\|_1 \right) \left(\frac{d^{\frac{L-i}{2}} s^{1-\frac{1}{p}}}{\frac{1}{p} - 1} \right) \right) \quad (3)$$

Pruning error in (3) is summed over each network layer. Considering layer i , the green factor is inherited from Theorem 2 with an added exponent due to a recursion over network layers after i . Similarly, the blue factor accounts for propagation of error through weight matrices after layer i , revealing that green and blue factors account for propagation of error through network layers. The red portion of (3), which captures the convergence properties of multi-layer pruning, comes from Lemma 3, where an extra factor $d^{\frac{L-i}{2}}$ arises as an artifact of the proof.⁸ Because d is fixed multiple of s determined by the pruning ratio, this factor disappears when p is small, leading the expression to converge. For example, if $p = \frac{1}{4}$, the red expression in (3) behaves asymptotically as $\{\mathcal{O}(s^{-2.5}), \mathcal{O}(s^{-2}), \mathcal{O}(s^{-1.5}), \dots\}$ for layers at the end of the network moving backwards.

6. Experiments

Within this section, we provide empirical analysis of i -SpaSP. We first present synthetic results using two-layer neural networks to numerically verify Theorem 4. Then, we perform experiments with two-layer networks on MNIST (Deng, 2012), convolutional neural networks (CNNs) on ImageNet (ILSVRC2012), and multi-lingual BERT (mBERT) (Devlin et al., 2018) on the cross-lingual NLI corpus (XNLI) (Conneau et al., 2018). For all experiments, we adopt best practices from previous work (Li et al., 2016) to determine pruning ratios within the dense network, often performing less pruning on sensitive layers; see Appendix ?? for details. As baselines, we adopt both greedy selection methodologies (Ye et al., 2020; Ye et al., 2020) and several common, heuristic methods. We find that, in addition to improving upon the pruning efficiency




Figure 2: i -SpaSP pruning experiments for two-layer networks of different sizes and p ratios. The error decay rate increases with decreasing p , confirming that more compressible hidden representations aid in the pruning process.

⁸. This artifact arises because ℓ_1 norms must be replaced with ℓ_2 norms in certain areas to form a recursion over network layers. This artifact can likely be removed by leveraging sparse matrix analysis for expander graphs (Bah and Tanner, 2018), but we leave this as future work to simplify the analysis.

of GFS (i.e., see Section 4.2 for runtime comparisons on ResNet pruning experiments), i-SpaSP performs comparably to baselines in all cases, demonstrating that it is both performant and efficient.

6.1. Synthetic Experiments

To numerically verify Theorem 4, we construct synthetic H matrices with different row-compressibility ratios p , denoted as $\mathcal{H} = \{H^{(i)}\}_{i=1}^K$. For each $p \in \{0.3, 0.5, 0.7, 0.9\}$, we randomly generate three unique entries within \mathcal{H} (i.e., $K = 12$) and present average performance across them all. We prune a randomly-initialized $W^{(1)}$ from size $d_{hid} \times d_{out}$ to sizes $s \times d_{out}$ for $s \in [d_{hid}]$. (i.e., each setting of s is a separate experiment) using $T = 20$.⁹ We test $d_{hid} \in \{100, 200\}$, and the same $W^{(1)}$ matrix is used for experiments with equal hidden dimension; see Appendix ?? for more details.

Results are displayed in Figure 2, where $\|V\|_F^2$ is shown to decay polynomially with respect to the number of neurons within the pruned network s . Furthermore, as predicted by Theorem 4, the decay rate of $\|V\|_F^2$ increases as p decreases, revealing that higher levels of sparsity within the dense network’s hidden representations improve pruning performance and speed with respect to s . This trend holds for all hidden dimensions that were considered.

6.2. Two-Layer Networks

We perform pruning experiments with two-layer networks on MNIST (Deng, 2012). All MNIST images are flattened and no data augmentation is used. The dense network has 10×10^3 hidden neurons and is pre-trained before pruning. We prune the dense network using i-SpaSP, GFS, and Top-K, which we use as a naive greedy selection baseline.¹⁰ After pruning, the network is fine-tuned using stochastic gradient descent (SGD) with momentum. Performance is reported as an average across three separate trails; see Appendix ?? for more details.

As shown in Figure 3-right, i-SpaSP outperforms other pruning methodologies after fine-tuning in all experimental settings. Networks obtained with GFS perform well without fine-tuning (i.e., Figure 3-left) because neurons are selected to minimize loss during the pruning process. Because i-SpaSP selects neurons based upon the importance criteria described in Section 4, the pruning process does not directly minimize training loss, thus leading to poorer performance prior to fine-tuning (i.e., Top-K exhibits similar behavior). Nonetheless, i-SpaSP, in addition to improving upon the pruning efficiency of GFS, discovers a set




Figure 3: Performance of networks pruned with different greedy algorithms on MNIST before (left) and after (right) fine-tuning. Although GFS performs well prior to fine-tuning, *i-SpaSP* always yields the top-performing network after fine-tuning.

9. We find that the active set of neurons selected by i-SpaSP becomes stable (i.e., few neurons are modified) after 20 iterations or less in all synthetic experiments.

10. Top-k selects k neurons with the largest-magnitude hidden representations in a mini-batch of data. It is a naive baseline for greedy selection that is not used in previous work to the best of our knowledge.

of neurons that more closely recovers dense network output, as revealed by its superior performance after fine-tuning.

6.3. Deep Convolutional Networks

We perform structured filter pruning of ResNet34 (He et al., 2015) and MobileNetV2 (Sandler et al., 2018) with i-SpaSP on ImageNet (ILSVRC 2012). Beginning with public, pre-trained models (Paszke et al., 2017), we use i-SpaSP to prune chosen convolutional blocks within each network, then fine-tune the pruned model using SGD with momentum. After pruning each block, we perform a small amount of fine-tuning; see Appendix ?? for more details. Numerous heuristic and greedy selection-based algorithms are adopted as baselines; see Table 1.

ResNet34. We prune ResNet34 to 2.69 and 2.13 GFlops with i-SpaSP. As shown in Table 1, i-SpaSP yields comparable performance to GFS variants at similar FLOP levels; e.g., i-SpaSP matches 75.5% test accuracy of GFS with the 2.69 GFlop model and performs within 1% of GFS variants for the 2.13 GFlop model. However, the multi-layer variant of GFS (Ye et al., 2020) does discover sub-networks with fewer FLOPS and similar performance in both cases. i-SpaSP improves upon or matches the performance of all heuristic methods at similar FLOP levels.

MobileNetV2. We prune MobileNetV2 to 260, 242, and 220 MFlops with i-SpaSP. In all cases, sub-networks discovered with i-SpaSP achieve performance within 1% of those obtained via both GFS variants and heuristic methods. Although MobileNetV2 performance is relatively lower in comparison to ResNet34, i-SpaSP is still capable of pruning the more difficult network to various different FLOP levels and performs comparably to baselines in all cases.

Discussion. Within Table 1, the performance of i-SpaSP is never more than 1% below that of a similar-FLOP model obtained with a baseline pruning methodology, including both greedy selection and heuristic-based methods. As such, similar to GFS, i-SpaSP can be seen as a theoretically-grounded pruning methodology that is practically useful, even in large-scale experiments. Such

	Pruning Method	Top-1 Acc.	FLOPS	
ResNet34	Full Model He et al. (2015)	73.4	3.68G	
	Filter Pruning Li et al. (2016)	72.1	2.79G	
	Rethinking Pruning Liu et al. (2018)	72.0	2.79G	
	More is Less Dong et al. (2017)	73.0	2.75G	
	i-SpaSP	73.5	2.69G	
	GFS Ye et al. (2020)	73.5	2.64G	
	Multi-Layer GFS Ye et al. (2020)	73.5	2.20G	
	SFP He et al. (2018a)	71.8	2.17G	
	FPGM He et al. (2019)	72.5	2.16G	
	i-SpaSP	72.5	2.13G	
	GFS Ye et al. (2020)	72.9	2.07G	
	Multi-Layer GFS Ye et al. (2020)	73.3	1.90G	
	MobileNetV2	Full Model Sandler et al. (2018)	72.0	314M
		i-SpaSP	71.6	260M
GFS Ye et al. (2020)		71.9	258M	
Multi-Layer GFS Ye et al. (2020)		72.2	245M	
i-SpaSP		71.3	242M	
LEGR Chin et al. (2019)		71.4	224M	
Uniform		70.0	220M	
AMC He et al. (2018b)		70.8	220M	
i-SpaSP		70.7	220M	
GFS Ye et al. (2020)		71.6	220M	
Multi-Layer GFS Ye et al. (2020)	71.7	218M		
Meta Pruning Liu et al. (2019)	71.2	217M		

Table 1: Test accuracy of ResNet34 and MobileNetV2 models pruned to different FLOP levels with various pruning algorithms on ImageNet.

pruning methodologies that are both theoretically and practically relevant are few. In comparison to GFS variants, i-SpaSP significantly improves pruning efficiency; see Section 4.2. Thus, i-SpaSP can be seen as a viable alternative to GFS—both theoretically and practically—that may be preferable when runtime is a major concern.

6.4. Transformer Networks

We perform structured pruning of attention heads using the mBERT model (Devlin et al., 2018) on the XNLI dataset (Conneau et al., 2018), which contains textual entailment annotations across 15 different languages. We begin with a pre-trained mBERT model and fine-tune it on XNLI prior to pruning. Then, structured pruning is performed on the attention heads of each layer using i-SpaSP, uniform pruning (i.e., randomly removing a fixed ratio of attention heads), and a sensitivity-based masking approach (Michel et al., 2019) (i.e., a state-of-the-art heuristic approach for structured attention head pruning for transformers). After pruning models to a fixed ratio of 25% or 40% of original attention heads, we fine-tune the pruned networks and record their performance on the XNLI test set; see Appendix ?? for further details.

As shown in Table 2, models pruned with i-SpaSP outperform those pruned to the same ratio with either uniform or heuristic pruning methods in all cases. The ability of i-SpaSP to effectively prune transformers demonstrates that the methodology can be applied to the structured pruning of numerous different architectures without significant implementation changes (i.e., using the automatic differentiation approach described in Section 4.1). Furthermore, i-SpaSP even outperforms a state-of-the-art heuristic approach for the structured pruning of transformer attention heads, thus again highlighting the strong empirical performance of i-SpaSP in comparison to heuristic pruning methods that lack theoretical guarantees.

7. Conclusion

We propose i-SpaSP, a pruning methodology for neural networks inspired by sparse signal recovery. Our methodology comes with theoretical guarantees that indicate, for both two and multi-layer networks, the quality of a pruned network decays polynomially with respect to its size. We connect this theoretical analysis to properties of the dense network, showing that pruning performance improves as the dense network’s hidden representations become more sparse. Practically, i-SpaSP performs comparably to numerous baseline pruning methodologies in large-scale experiments and drastically improves upon the computational efficiency of the most common provable pruning methodologies. As such, i-SpaSP is a practical, provable, and efficient algorithm that we hope will enable a better understanding of neural network pruning both in theory and practice. In future work, we wish to extend i-SpaSP to cover cases in which network pruning and training are combined into a single process, such as utilizing regularization-based approaches to induce sparsity during pre-training or updating network weights to improve sub-network performance as pruning occurs.

Method	Pruning Ratio	Top-1 Acc.
Full Model	-	71.02
Uniform	25%	62.73
	40%	66.43
Michel et al. (2019)	25%	62.97
	40%	67.31
i-SpaSP	25%	63.70
	40%	68.11

Table 2: Test accuracy of mBERT models pruned and fine-tuned on the XNLI dataset.

Acknowledgments

Compute resources used for this work were provided by Alegion Inc. This work is supported by NSF FET:Small no. 1907936, NSF MLWiNS CNS no. 2003137 (in collaboration with Intel), NSF CMMI no. 2037545, NSF CAREER award no. 2145629, and Rice InterDisciplinary Excellence Award (IDEA).

References

- Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dan Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015.
- Francis Bach. Breaking the Curse of Dimensionality with Convex Neural Networks. *arXiv e-prints*, art. arXiv:1412.8690, December 2014.
- Bubacarr Bah and Jared Tanner. On the construction of sparse matrices from expander graphs. *Frontiers in Applied Mathematics and Statistics*, 4:39, 2018.
- Cenk Baykal, Lucas Liebenwein, Igor Gilitschenski, Dan Feldman, and Daniela Rus. Data-dependent coresets for compressing neural networks with applications to generalization bounds. *arXiv preprint arXiv:1804.05345*, 2018.
- Cenk Baykal, Lucas Liebenwein, Igor Gilitschenski, Dan Feldman, and Daniela Rus. SiPPing Neural Networks: Sensitivity-informed Provable Pruning of Neural Networks. *arXiv e-prints*, art. arXiv:1910.05422, October 2019.
- Emmanuel J Candes and Terence Tao. Near-optimal signal recovery from random projections: Universal encoding strategies? *IEEE transactions on information theory*, 52(12):5406–5425, 2006.
- Xiaohan Chen, Yu Cheng, Shuohang Wang, Zhe Gan, Zhangyang Wang, and Jingjing Liu. EarlyBERT: Efficient BERT Training via Early-bird Lottery Tickets. *arXiv e-prints*, art. arXiv:2101.00063, December 2020.
- Ting-Wu Chin, Ruizhou Ding, Cha Zhang, and Diana Marculescu. Legr: Filter pruning via learned global ranking. 2019.
- Alexis Conneau, Ruty Rinott, Guillaume Lample, Adina Williams, Samuel R. Bowman, Holger Schwenk, and Veselin Stoyanov. Xnli: Evaluating cross-lingual sentence representations. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, 2018.

- Li Deng. The mnist database of handwritten digit images for machine learning research. *IEEE Signal Processing Magazine*, 29(6):141–142, 2012.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- Xuanyi Dong, Junshi Huang, Yi Yang, and Shuicheng Yan. More is less: A more complicated network with less inference complexity. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5840–5848, 2017.
- Alan Edelman and Yuyang Wang. Random matrix theory and its innovative applications. In *Advances in Applied Mathematics, Modeling, and Computational Science*, pages 91–116. Springer, 2013.
- Utku Evci, Trevor Gale, Jacob Menick, Pablo Samuel Castro, and Erich Elsen. Rigging the Lottery: Making All Tickets Winners. *arXiv e-prints*, art. arXiv:1911.11134, November 2019.
- Utku Evci, Yani A. Ioannou, Cem Keskin, and Yann Dauphin. Gradient Flow in Sparse Neural Networks and How Lottery Tickets Win. *arXiv e-prints*, art. arXiv:2010.03533, October 2020.
- Marguerite Frank, Philip Wolfe, et al. An algorithm for quadratic programming. *Naval research logistics quarterly*, 3(1-2):95–110, 1956.
- Jonathan Frankle and Michael Carbin. The Lottery Ticket Hypothesis: Finding Sparse, Trainable Neural Networks. *arXiv e-prints*, art. arXiv:1803.03635, March 2018.
- Jonathan Frankle, Gintare Karolina Dziugaite, Daniel M. Roy, and Michael Carbin. Stabilizing the Lottery Ticket Hypothesis. *arXiv e-prints*, art. arXiv:1903.01611, March 2019.
- Song Han, Huizi Mao, and William J. Dally. Deep Compression: Compressing Deep Neural Networks with Pruning, Trained Quantization and Huffman Coding. *arXiv e-prints*, art. arXiv:1510.00149, October 2015.
- Song Han, Huizi Mao, and William J Dally. Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding. *arXiv preprint arXiv:1510.00149*, 2015a.
- Song Han, Jeff Pool, John Tran, and William J Dally. Learning both weights and connections for efficient neural networks. *arXiv preprint arXiv:1506.02626*, 2015b.
- Song Han, Xingyu Liu, Huizi Mao, Jing Pu, Ardavan Pedram, Mark A. Horowitz, and William J. Dally. EIE: Efficient Inference Engine on Compressed Deep Neural Network. *arXiv e-prints*, art. arXiv:1602.01528, February 2016.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep Residual Learning for Image Recognition. *arXiv e-prints*, art. arXiv:1512.03385, December 2015.
- Yang He, Guoliang Kang, Xuanyi Dong, Yanwei Fu, and Yi Yang. Soft filter pruning for accelerating deep convolutional neural networks. *arXiv preprint arXiv:1808.06866*, 2018a.
- Yang He, Ping Liu, Ziwei Wang, Zhilan Hu, and Yi Yang. Filter pruning via geometric median for deep convolutional neural networks acceleration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4340–4349, 2019.

- Yihui He, Xiangyu Zhang, and Jian Sun. Channel Pruning for Accelerating Very Deep Neural Networks. *arXiv e-prints*, art. arXiv:1707.06168, July 2017.
- Yihui He, Ji Lin, Zhijian Liu, Hanrui Wang, Li-Jia Li, and Song Han. Amc: Automl for model compression and acceleration on mobile devices. In *Proceedings of the European conference on computer vision (ECCV)*, pages 784–800, 2018b.
- Zehao Huang and Naiyan Wang. Data-driven sparse structure selection for deep neural networks. In *Proceedings of the European conference on computer vision (ECCV)*, pages 304–320, 2018.
- Rajiv Khanna and Anastasios Kyriillidis. It dies hard: Provable accelerated iterative hard thresholding. In *International Conference on Artificial Intelligence and Statistics*, pages 188–198. PMLR, 2018.
- Hao Li, Asim Kadav, Igor Durdanovic, Hanan Samet, and Hans Peter Graf. Pruning Filters for Efficient ConvNets. *arXiv e-prints*, art. arXiv:1608.08710, August 2016.
- Lucas Liebenwein, Cenk Baykal, Harry Lang, Dan Feldman, and Daniela Rus. Provable filter pruning for efficient neural networks. *arXiv preprint arXiv:1911.07412*, 2019.
- Zechun Liu, Haoyuan Mu, Xiangyu Zhang, Zichao Guo, Xin Yang, Kwang-Ting Cheng, and Jian Sun. Metapruning: Meta learning for automatic neural network channel pruning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3296–3305, 2019.
- Zhuang Liu, Jianguo Li, Zhiqiang Shen, Gao Huang, Shoumeng Yan, and Changshui Zhang. Learning Efficient Convolutional Networks through Network Slimming. *arXiv e-prints*, art. arXiv:1708.06519, August 2017.
- Zhuang Liu, Mingjie Sun, Tinghui Zhou, Gao Huang, and Trevor Darrell. Rethinking the Value of Network Pruning. *arXiv e-prints*, art. arXiv:1810.05270, October 2018.
- Jian-Hao Luo, Jianxin Wu, and Weiyao Lin. ThiNet: A Filter Level Pruning Method for Deep Neural Network Compression. *arXiv e-prints*, art. arXiv:1707.06342, July 2017.
- Eran Malach, Gilad Yehudai, Shai Shalev-Shwartz, and Ohad Shamir. Proving the Lottery Ticket Hypothesis: Pruning is All You Need. *arXiv e-prints*, art. arXiv:2002.00585, February 2020.
- Paul Michel, Omer Levy, and Graham Neubig. Are sixteen heads really better than one? *Advances in neural information processing systems*, 32, 2019.
- Pavlo Molchanov, Stephen Tyree, Tero Karras, Timo Aila, and Jan Kautz. Pruning convolutional neural networks for resource efficient inference. *arXiv preprint arXiv:1611.06440*, 2016.
- Ben Mussay, Margarita Osadchy, Vladimir Braverman, Samson Zhou, and Dan Feldman. Data-independent neural pruning via coresets. *arXiv preprint arXiv:1907.04018*, 2019.
- Deanna Needell and Joel A Tropp. Cosamp: Iterative signal recovery from incomplete and inaccurate samples. *Applied and computational harmonic analysis*, 26(3):301–321, 2009.
- George L Nemhauser, Laurence A Wolsey, and Marshall L Fisher. An analysis of approximations for maximizing submodular set functions—i. *Mathematical programming*, 14(1):265–294, 1978.

- Laurent Orseau, Marcus Hutter, and Omar Rivasplata. Logarithmic Pruning is All You Need. *arXiv e-prints*, art. arXiv:2006.12156, June 2020.
- Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. 2017.
- Ankit Pensia, Shashank Rajput, Alliot Nagle, Harit Vishwakarma, and Dimitris Papailiopoulos. Optimal lottery tickets via subsetsum: Logarithmic over-parameterization is sufficient. *arXiv preprint arXiv:2006.07990*, 2020.
- Sebastian Pokutta, Christoph Spiegel, and Max Zimmer. Deep Neural Network Training with Frank-Wolfe. *arXiv e-prints*, art. arXiv:2010.07243, October 2020.
- Vivek Ramanujan, Mitchell Wortsman, Aniruddha Kembhavi, Ali Farhadi, and Mohammad Rastegari. What’s Hidden in a Randomly Weighted Neural Network? *arXiv e-prints*, art. arXiv:1911.13299, November 2019.
- Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4510–4520, 2018.
- Xavier Suau, Nicholas Apostoloff, et al. Filter distillation for network compression. In *2020 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 3129–3138. IEEE, 2020.
- Chaoqi Wang, Guodong Zhang, and Roger Grosse. Picking Winning Tickets Before Training by Preserving Gradient Flow. *arXiv e-prints*, art. arXiv:2002.07376, February 2020.
- Cameron R Wolfe, Qihan Wang, Junhyung Lyle Kim, and Anastasios Kyrillidis. Provably efficient lottery ticket discovery. *arXiv preprint arXiv:2108.00259*, 2021.
- Mao Ye. Network-pruning-greedy-forward-selection. <https://github.com/lushleaf/Network-Pruning-Greedy-Forward-Selection>, 2021.
- Mao Ye, Chengyue Gong, Lizhen Nie, Denny Zhou, Adam Klivans, and Qiang Liu. Good sub-networks provably exist: Pruning via greedy forward selection. In *International Conference on Machine Learning*, pages 10820–10830. PMLR, 2020.
- Mao Ye, Lemeng Wu, and Qiang Liu. Greedy Optimization Provably Wins the Lottery: Logarithmic Number of Winning Tickets is Enough. *arXiv e-prints*, art. arXiv:2010.15969, October 2020.
- Haoran You, Chaojian Li, Pengfei Xu, Yonggan Fu, Yue Wang, Xiaohan Chen, Richard G Baraniuk, Zhangyang Wang, and Yingyan Lin. Drawing early-bird tickets: Towards more efficient training of deep networks. *arXiv preprint arXiv:1909.11957*, 2019.
- Ruichi Yu, Ang Li, Chun-Fu Chen, Jui-Hsin Lai, Vlad I. Morariu, Xintong Han, Mingfei Gao, Ching-Yung Lin, and Larry S. Davis. NISP: Pruning Networks using Neuron Importance Score Propagation. *arXiv e-prints*, art. arXiv:1711.05908, November 2017.

Shuai Zhang, Meng Wang, Sijia Liu, Pin-Yu Chen, and Jinjun Xiong. Why lottery ticket wins? a theoretical perspective of sample complexity on pruned neural networks. *arXiv preprint arXiv:2110.05667*, 2021.

Zhuangwei Zhuang, Mingkui Tan, Bohan Zhuang, Jing Liu, Yong Guo, Qingyao Wu, Junzhou Huang, and Jinhui Zhu. Discrimination-aware Channel Pruning for Deep Neural Networks. *arXiv e-prints*, art. arXiv:1810.11809, October 2018.