

# LEARNING STRUCTURED SPARSE NEURAL NETWORKS USING GROUP ENVELOPE REGULARIZATION

**Anonymous authors**

Paper under double-blind review

## ABSTRACT

We propose an efficient method to learn both unstructured and structured sparse neural networks during training, utilizing a novel generalization of the sparse envelope function (SEF) used as a regularizer, termed *weighted group sparse envelope function* (WGSEF). The WGSEF acts as a neuron group selector, which is leveraged to induce structured sparsity. The method ensures a hardware-friendly structured sparsity of a deep neural network (DNN) to efficiently accelerate the DNN’s evaluation. Notably, the method is adaptable, letting any hardware specify group definitions, such as filters, channels, filter shapes, layer depths, a single parameter (unstructured), etc. Owing to the WGSEF’s properties, the proposed method allows to a pre-define sparsity level that would be achieved at the training convergence, while maintaining negligible network accuracy degradation or even improvement in the case of redundant parameters. We introduce an efficient technique to calculate the exact value of the WGSEF along with its proximal operator in a worst-case complexity of  $O(n)$ , where  $n$  is the total number of group variables. In addition, we propose a proximal-gradient-based optimization method to train the model, that is, the non-convex minimization of the sum of the neural network loss and the WGSEF. Finally, we conduct an experiment and illustrate the efficiency of our proposed technique in terms of the completion ratio, accuracy, and inference latency.

## 1 INTRODUCTION

In the last decade, far-reaching progress has been achieved in the study of neural networks, as these achieve the best performance in many machine learning tasks. Nonetheless, modern networks are becoming gradually larger with an enormous number of parameters with the cost of rising storage, memory footprint, computing resources and energy consumption Deng et al. (2020); Cheng et al. (2017). Notwithstanding, it has been shown in recent studies Han et al. (2015a); Ullrich et al. (2017); Molchanov et al. (2017) that the latest neural networks tend to be highly over-parametrized in the sense that a relatively large amount of parameters are redundant and could have been pruned (or, nullified) without degrading the network precision.

Over-parametrization has two notable problematic effects: first, it can very easily results in overfitting Allen-Zhu et al. (2019), as well as memorizing random patterns in the data Zhang et al. (2021), leading to inferior generalization. Second, it becomes quite challenging to deploy these networks to low hardware edge devices due to, e.g., high computational cost results in a large consumption of power and a long inference time, and a substantial redundant memory space for storing is needed (which obviously would never be available on edge devices). This has led to the search for sparse efficient and effective neural network architectures which has become increasingly a front challenge.

The most common way to tackle the first issue, namely, avoiding overfitting for better generalization is, to use properly suitable regularization, while the accepted approach to address the second problem is to apply post-training model compression techniques. However, the sparsification technique, i.e., regularizing the number of parameters during training, would result in both model compression and overfit-regularization.

Quite recently, the idea of structured sparsification was used in Wen et al. (2016b); Bui et al. (2021) to learn sparse neural networks that leverage tensor arithmetic in dedicated neural processing units

(NPU). In a nutshell, structured sparsity learning amounts to inducing sparsity onto structured components (e.g., channels, filters, or layers) in the neural network during the optimization procedure. This leads, in practice, to both low latency and lower power consumption, which can not be obtained by deploying unstructured sparse models on such modern hardware.

In the case of the unstructured sparsity-inducing, the most natural regularizer would be the so-called  $\ell_0$ -pseudo-norm function that counts the number of nonzero elements in the input vector, i.e.,  $\|\mathbf{z}\|_0 \triangleq |\{i : z_i \neq 0\}|$ . These sparse regularized minimization/ training problems are of the form  $\min_{\mathbf{z} \in \mathbb{R}^n} \{f(\mathbf{z}) + \lambda \|\mathbf{z}\|_0\}$ , or, alternatively, one can explicitly constrain the number of parameters used for regression and solve  $\min_{\mathbf{z} \in \mathbb{R}^n} \{f(\mathbf{z}) : \|\mathbf{z}\|_0 \leq k\}$ . Unfortunately, the  $\ell_0$ -norm is a difficult function to handle being non-convex and even non-continuous. Indeed, these types of regression models are known to be NP-hard problems in general Natarajan (1995). The only known algorithm guaranteed to output the global optimal solution is essentially a brute-force search over all possible subsets components of the vector  $\mathbf{z}$ . In practice, due to this expensive computational cost, the global optimal solution can not be computed in a reasonable time, even for a very small number of parameters, precluding its use in large neural networks.

As a remedy for the inherent problem above, Beck & Refael (2022) proposed a highly efficient tractable tight convex relaxation technique, termed sparse envelope function (SEF), for the sum of both  $\ell_0$  and  $\ell_2$  norms. Specifically, Beck & Refael (2022) suggested using this relaxation as a regularizer term for a convex loss objective, particularly for a linear regression model, to achieve feature selection while explicitly limiting the number of features to be a fixed parameter  $k$ . It was also shown that the performance of the regularization method in both reconstruction of a sparse noisy signal and recovering its support, surpass the performance of state-of-the-art techniques, such as, the Elastic-net Zou & Hastie (2005),  $k$ -support norm Argyriou et al. (2012), etc. Also, it was shown that the computational complexity of the SEF approach is linear in the number of parameters, while all others requires at least quadratic in the number of features, and thus very attractive.

With the goal of enabling structured sparsification learning that can be customized for different NPU devices, in this paper we propose a novel generalized notion of the SEF regularizer to handle group structured sparsification in neural network training. Our new generalized regularization term selects the most essential  $k \leq m$  predefined groups of neurons (which could be convolutional filters, channels, individual neurons, or any other user-defined/NPU definition) and prunes all others, while maintaining minimal network accuracy degradation. We define the new regularization term mathematically, propose an efficient method to calculate its value and proximal operator, and suggest a new algorithm to solve the complete optimization problem involving the non-convex term, which is the composition of the loss function and the neural network output.

**Related work.** The topic of regularization-based pruning received a lot of attention in recent years. Generally speaking, these studies can be divided into unstructured and structured pruning. Most prominent regularizers are the convex  $\ell_1$  and  $\ell_2$  norms Liu et al. (2017); Ye et al. (2018); Han et al. (2015a;b), as well as the non-convex  $\ell_0$  “norm” Louizos et al. (2017); Han et al. (2015c), where Bayesian methods and additional regularization terms for practicality, were used to deal with the non-convexity of the  $\ell_0$  norm. Additional works of Donoho & Elad (2003); com (2021-2023) suggest methods for norm  $\ell_0$  relaxation by employing  $l_1$  minimization in general (nonorthogonal) dictionaries and leading to an error surface with fewer local minima than the  $l_0$  norm. The motivation for these regularizers is their “sparsity-inducing” property which can be harnessed to learn sparse neural networks. While these fundamental papers significantly reduce the storage needed to store the networks on hardware, There were no benefits in reducing the inference latency time or either in cutting down power consumption. That is, the sparse neural networks, learned by the aforementioned methods, were not adapted to the structure of the hardware they aimed to run on. This observation led researchers to propose regularization-based structured pruning in favor of accelerating the running time. For example, Lebedev & Lempitsky (2015); Wen et al. (2016a); Yuan & Lin (2006) proposed the use of the Group Lasso regularization technique to learn sparse structures, and Scardapane et al. (2017) uses Sparse Group Lasso, summing Group Lasso with the standard Lasso penalty. Other convex regularizers include the Combined Group and Exclusive Sparsity (CGES) Yoon & Hwang (2017), which extends Exclusive Lasso (in essence, squared  $\ell_1$  over groups) Zhou et al. (2010) using Group Lasso. Recently, Bui et al. (2021) suggested a family of nonconvex regularizers that blend Group Lasso with nonconvex terms ( $\ell_0$ ,  $\ell_1 - \ell_2$  Lou et al. (2015), and SCAD Fan & Li (2001)). Since Bui et al. (2021) introduces non-convexity term into the penalty, it also requires

an appropriate optimization scheme, for which the authors propose an Augmented Lagrangian type method. However, this optimization algorithm has an inner optimization loop with a high computational cost. Moreover, their extensive experiments do not show an accuracy or sparsity advantages over convex penalties, suggesting that it might be still desirable to use a convex regularizer. Other methods, such as Chen et al. (2021); Li et al. (2019), focus on a group structure that captures the relations between parameters, neurons, and layers, in order to construct groups that can maximize network compression while minimizing accuracy loss. However, these methods still apply Group Lasso regularization. Specifically in Chen et al. (2021), the authors introduce the concept of Zero-Invariant Groups (ZIGs), which includes all input and output connections between layers. In the context of CNNs, it extends the channel-wise grouping Wen et al. (2016b) to include corresponding batch normalization parameters. By using this group structure, entire blocks of parameters can be removed while keeping dimensions aligned between layers, and ultimately allowing network compression. Moreover, their optimization scheme utilizes a two-phase algorithm to include a half-space projection step, which they name HSPG.

Finally, we mention that there exist other techniques for neural net compression, such as, quantization, low-rank decomposition, to name a few. In quantization, Courbariaux et al. (2016); Rastegari et al. (2016); Gong et al. (2014), the precision of the weights is reduced, by representing weights using a low number of bits (i.e., 8-bit) instead of higher one (i.e., 32-bit floating point values). The low-rank decomposition approach Denton et al. (2014); Jaderberg et al. (2014); Lebedev et al. (2014) is based on the observation that many weight matrices in neural networks are highly correlated and can be well approximated by matrices with a lower rank. By decomposing a weight matrix into lower-rank matrices, one can reduce the total number of parameters in the network.

**Notation.** We denote  $\mathbf{e}$  for the vector of all ones. For a positive integer  $m$ , we denote  $[m] \equiv \{1, 2, \dots, m\}$ . Let  $f : \mathbb{R}^n \rightarrow \mathbb{R}$ , be an extended real-valued function. Then, the conjugate function of  $f$ , denoted by  $f^* : \mathbb{R}^n \rightarrow \mathbb{R}$  is defined as  $f^*(y) = \max_{x \in \mathbb{R}^n} \{\langle x, y \rangle - f(x)\}$ , for any  $y \in \mathbb{R}^n$ . The bi-conjugate function is defined as the conjugate of the conjugate function, i.e.,  $f^{**}(x) = \max_{y \in \mathbb{R}^n} \{\langle x, y \rangle - f^*(y)\}$ , for any  $x \in \mathbb{R}^n$ . Finally, the proximal operator of a proper, lower semi-continuous convex function  $f : \mathbb{R}^n \rightarrow \mathbb{R}$  is defined as  $\text{prox}_f(v) = \arg \min_{x \in \mathbb{R}^n} \{f(x) + \frac{1}{2} \|x - v\|_2^2\}$ , for any  $v \in \mathbb{R}^n$ .

## 2 PROBLEM FORMULATION

In this section, we formulate the problem and introduce the weighted group sparse envelop function. Without loss of generality, our method is formulated on weights sparsity, but it can be directly extended to neuron sparsity (i.e., both weights and bias). Let  $\mathcal{D}$  be a dataset consisting of  $N$  i.i.d. input output pairs  $\{(x_1, y_1), \dots, (x_N, y_N)\}$ . The neural network training problem is formalized as the following empirical risk minimization procedure on the parameters  $\theta \in \Theta$  of a hypothesis  $f(\cdot; \theta)$ ,

$$\arg \min_{\theta \in \Theta} \frac{1}{N} \sum_{i=1}^N \mathcal{L}(f(x_i, \theta), y_i), \quad (1)$$

where  $\mathcal{L}(\cdot)$  corresponds to a loss function, e.g., cross-entropy loss for classification, mean-squared error for regression, etc. Below,  $n \triangleq |\theta|$  denotes the number of parameters. The training of neural networks is prone to overfitting, therefore, regularization is typically employed. Consider a regularized empirical risk minimization procedure with a regularizer  $\Omega(\cdot) : \Theta \rightarrow \mathbb{R}$  on the parameters  $\theta$  of a hypothesis  $f(\cdot; \theta)$ ,

$$\arg \min_{\theta \in \Theta} \frac{1}{N} \sum_{i=1}^N \mathcal{L}(f(x_i, \theta), y_i) + \lambda \cdot \Omega(\theta), \quad (2)$$

where  $\lambda \in \mathbb{R}$ . Weight-decay, known as the  $\ell_2$ -norm regularization, is the most common technique for deep network training regularization (in this case,  $\Omega(\theta) = \frac{1}{2} \|\theta\|_2^2$ ). It prevents overfitting and improves generalization as it enforces the weights to decrease proportionally to their magnitudes.

The most natural way to force a predefined  $k$ -level sparsity would be to constraint the number of non-zeros parameters (e.g., the model weights), which can be done by adding the constraint that

$\|\theta\|_0 \leq k$ , where  $k \leq n$  is the required predefined level of sparsity. In this case, the training problem is formalized as follows,

$$\operatorname{argmin}_{\theta \in \Theta} \frac{1}{N} \sum_{i=1}^N \mathcal{L}(f(x_i, \theta), y_i) + \frac{\lambda}{2} \|\theta\|_2^2 \quad \text{s.t.} \quad \|\theta\|_0 \leq k. \quad (3)$$

We refer to the above training problem as *unstructured sparsity*. In the case of *structured sparsity*, the parameters  $\theta$  are divided into *predefined* disjoint sub-groups. These subgroups could define building blocks architecture of DNN's, i.e., filters, channels, filter shapes, and layer depth. Consider the following definition.

**Definition 1** (Projection). *Let  $s$  be a subset of indexes  $s \subseteq \{1, 2, \dots, n\}$  of size  $|s| \leq n$ . Then, given some vector  $\theta \in \mathbb{R}^n$ , the projection  $M_s : \mathbb{R}^n \rightarrow \mathbb{R}^n$  preserves only the entries of  $\theta$  that belong to the set  $s$ . Furthermore, let  $A_s$  be an  $n \times n$  diagonal matrix, where  $[A_s]_{ii} = 1$  if  $i \in s$ , and zero, otherwise. Note that  $M_s(\theta) = A_s \theta$ .*

**Example 1.** *Let  $n = 3, \theta = (3, 6, 9)^\top, s = \{1, 3\} \subseteq [n]$ , and accordingly  $|s| = 2$ , then  $M_s(\theta) = M_s((3, 6, 9)^\top) = (3, 0, 9)^\top$ , with  $[A_s]_{11} = [A_s]_{33} = 1$ , and zero otherwise.*

Following the above definition, let  $m \leq n$  subsets  $s_1, s_2, \dots, s_m$  be a given partition of  $[n]$ , namely,  $s_i \cap s_j = \emptyset$ , for all  $i \neq j$ , and  $\bigcup_{i=1}^m s_i = [n]$ . Without loss of generality, we assume that  $n \% m = 0$ ; otherwise, the groups would have different coordinates. Every group is associated with some weight  $d_j \in \mathbf{R}_{++}$ , where  $j \in [m]$ , e.g.,  $d_j = \frac{1}{|s_j|}$ , namely, we normalize by the group size. For simplicity of notation, let  $\theta_{s_i} = M_{s_i}(\theta)$ , for  $i = 1, 2, \dots, m$ . Then, our structured training problem is,

$$\operatorname{argmin}_{\theta \in \Theta} \frac{1}{N} \sum_{i=1}^N \mathcal{L}(f(x_i, \theta), y_i) + \frac{\lambda}{2} \sum_{j=1}^m d_j \|\theta_{s_j}\|_2^2 \quad (4)$$

s.t.  $\|\|\theta_{s_1}\|_2^2, \|\theta_{s_2}\|_2^2, \dots, \|\theta_{s_m}\|_2^2\|_0 \leq k$ .

To wit, we constrain the number of groups which has at least one non-zero coordinate, to be at most  $k$ . Let  $C_k$  denote the set of all  $k$  sparse groups, i.e.,

$$C_k \triangleq \{\theta : \|\|\theta_{s_1}\|_2^2, \|\theta_{s_2}\|_2^2, \dots, \|\theta_{s_m}\|_2^2\|_0 \leq k\}, \quad (5)$$

and define  $\delta_{C_k}$  as the following extended real-valued function,

$$\delta_{C_k}(\Theta) \triangleq \begin{cases} 0, & \|\|\Theta_{s_1}\|_2^2, \|\Theta_{s_2}\|_2^2, \dots, \|\Theta_{s_m}\|_2^2\|_0 \leq k, \\ \infty, & \text{else.} \end{cases} \quad (6)$$

Then, the optimization problem in equation 4 can be reformulated as,

$$\operatorname{argmin}_{\theta \in \Theta} \frac{1}{N} \sum_{i=1}^N \mathcal{L}(f(x_i, \theta), y_i) + \lambda \cdot g_{s_k}(\theta), \quad (7)$$

where  $g_{s_k}(\theta) \triangleq \frac{1}{2} \sum_{j=1}^m d_j \|\theta_{s_j}\|_2^2 + \delta_{C_k}(\Theta)$ . Equivalently,  $g_{s_k}(\theta)$  can be rewritten as  $g_{s_k}(\theta) \triangleq \frac{1}{2} \sum_{i=1}^n d_i \cdot \theta_i^2 + \delta_{C_k}(\Theta)$ , where  $d_i = d_j$  for every  $i \in s_j$ .

The  $\ell_0$ -norm that appears in 7, is a difficult function to handle being nonconvex and even non-continuous, making the problem an untraceable combinatorial NP hard problem Natarajan (1995). Following Beck & Refael (2022), one approach to deal with this inherent difficulty is to consider the best convex estimator of  $g_{s_k}(\cdot)$ . The later is its biconjgate function, namely,  $\mathcal{G}\mathcal{S}_k(\theta) = g_{s_k}^{**}(\theta)$ , which we refer to as the *weighted group sparse envelope function (WGSEF)*.

**Remark 1.** *Consider the case  $m = n$ , namely, every subset  $s_i, i \in [m]$  is a singleton and  $\forall j \in [m], d_j = 1$ . Here,  $g_{s_k}(\theta) = s_k(\theta)$ , where  $s_k(\theta) = \frac{1}{2} \|\theta\|_2^2$  if  $\|\theta\|_0 \leq k$  and  $s_k(\theta) = \infty$ , otherwise. Accordingly, in this case,  $\mathcal{G}\mathcal{S}_k(\theta) = \mathcal{S}_k(\theta) = s_k^{**}(\theta)$ , and  $s_k^{**}(\theta)$  is exactly the classical SEF, namely,  $\mathcal{S}_k(\cdot)$ . Therefore,  $g_{s_k}(\cdot)$  is indeed a new generalization of SEF to handle group sparsity.*

Thus, the path taken in this paper is to consider the following relaxation training/learning problem,

$$\operatorname{argmin}_{\theta \in \Theta} \frac{1}{N} \sum_{i=1}^N \mathcal{L}(f(x_i, \theta), y_i) + \mathcal{G}\mathcal{S}_k(\theta). \quad (8)$$

In the following section, we develop an efficient algorithm to calculate the value and the prox-operator Beck (2017) of WGSEF; these will play an essential ingredient when solving (8).

### 3 TIGHT CONVEX RELAXATION

Let us start by introducing some notation. For any  $\theta \in \mathbb{R}^n$  and  $m$  subgroups of indexes  $s_1, s_2, \dots, s_m \in [n]$ , we denote by  $M_{\langle s_i \rangle}(\theta)$  the corresponding subgroup of coordinates in  $\theta$  with the  $i$ th largest  $\ell_2$ -norm, i.e.,

$$\|M_{\langle s_1 \rangle}(\theta)\|_2 \geq \|M_{\langle s_2 \rangle}(\theta)\|_2 \geq \dots \geq \|M_{\langle s_m \rangle}(\theta)\|_2, \quad \forall i \neq j \in [m].$$

We next show that the conjugate of the  $k$  group sparse envelopes is the  $k$  weighted group hard thresholding function. In the sequel, we let  $D$  be the  $n \times n$  diagonal positive-definite weights matrix, such that  $D_{i,i} = \sqrt{d_i}, \forall i \in s_j$ , and  $j \in [m]$ .

**Lemma 3.1** (The  $k$  weighted group sparse envelop conjugate). *Let subsets  $s_1, s_2, \dots, s_m$  be a set of  $m \leq n$  disjoint indexes that partition  $[n]$ . Then, for any  $\tilde{\theta} \in \mathbb{R}^n$ ,*

$$gS_k^*(\tilde{\theta}) = \frac{1}{2} \sum_{j=1}^k \frac{1}{d_j} \left\| M_{\langle s_j \rangle}(\tilde{\theta}) \right\|_2^2. \quad (9)$$

Next, we obtain the bi-conjugate function of the  $k$  weighted group sparse envelope.

**Lemma 3.2** (The variational bi-conjugate  $k$  weighted group sparse envelop). *Let  $s_1, s_2, \dots, s_m$  be a set of  $m \leq n$  disjoint subsets that partition  $[n]$ . Then, for any  $\theta \in \mathbb{R}^n$ , the bi-conjugate of the  $k$  group sparse envelop is given by*

$$\mathcal{GS}_k(\theta) = \frac{1}{2} \min_{\mathbf{u} \in B_k} \left\{ \sum_{j=1}^m d_j \phi(A_{s_j} \theta, u_j) \right\}, \quad (10)$$

where

$$\phi(A_{s_j} \theta, u_j) \triangleq \begin{cases} \frac{\theta^\top A_{s_j} \theta}{u_j}, & u_j > 0, \\ 0, & u_j = 0 \cap A_{s_j} \theta = 0, \\ \infty & \text{else.} \end{cases} \quad (11)$$

The following is a straightforward corollary of Lemma 3.2.

**Corollary 3.2.1.** *The following holds:*

$$\mathcal{GS}_k(\theta) = \mathcal{S}((\sqrt{d_1} \|A_{s_1} \theta\|_2, \sqrt{d_2} \|A_{s_2} \theta\|_2, \dots, \sqrt{d_m} \|A_{s_m} \theta\|_2)^\top), \quad (12)$$

where  $\mathcal{S}(\theta) \triangleq s_k^{**}(\theta)$  is the standard SEF.

The implication of the above corollary is that in order to calculate  $\mathcal{GS}_k(\theta)$  we only need to apply an algorithm that calculates the SEF at  $(\sqrt{d_1} \|A_{s_1} \theta\|_2, \sqrt{d_2} \|A_{s_2} \theta\|_2, \dots, \sqrt{d_m} \|A_{s_m} \theta\|_2)^\top$ .

**Remark 2.** *Noting  $\|\sqrt{d_j} A_{s_j} \theta\|_2^2 = \sum_{i \in s_j} d_j \theta_i^2$  we observe that since  $s_j$  is given, the number of operations required to calculate  $\|\sqrt{d_j} A_{s_j} \theta\|_2^2$  is linear w.r.t.  $|s_j|$ . Thus, the computational complexity of calculating the vector  $(\sqrt{d_1} \|A_{s_1} \theta\|_2, \sqrt{d_2} \|A_{s_2} \theta\|_2, \dots, \sqrt{d_m} \|A_{s_m} \theta\|_2)^\top$  is linear in  $n$ .*

### 4 PROXIMAL MAPPING OF THE WGSEF

In this section, we will show how to efficiently compute the proximal operator of positive scalar multiples of  $\mathcal{GS}_k$ . The ability to perform such an operation implies that it is possible to employ fast proximal gradient methods to solve equation 8. We begin with the following lemma that shows that the proximal operator can be determined in terms of the optimal solution of a convex problem that resembles the optimization problem defined in Lemma equation 10 for computing  $\mathcal{GS}_k$ .

**Lemma 4.1.** *Let  $\lambda > 0$ ,  $t \in \mathbb{R}^n$ , and  $s_1, s_2, \dots, s_m$  be a set of  $m \leq n$  disjoint subsets that partition  $[n]$ . Then,  $v = \text{prox}_{\lambda \mathcal{GS}_k}(t)$  is given by*

$$j \in [m], \quad A_{s_j} v = \frac{u_j A_{s_j} t}{\lambda d_j + u_j}, \quad (13)$$

where  $(u_1, u_2, \dots, u_n)^T$  is the minimizer of

$$\min_{\mathbf{u} \in D_k} \sum_{j=1}^m \phi \left( \sqrt{d_j} A_{s_j} t, \lambda d_j + u_j \right). \quad (14)$$

Next, we show that the proximal operator of  $\mathcal{GS}_k$  reduces to an efficient one-dimensional search.

**Corollary 4.1.1** (The proximal operator of  $\mathcal{GS}_k$ ). *The solution  $u_j = u_j(\mu^*)$  of equation 14 with  $u_j(\cdot)$  defined as<sup>1</sup>*

$$u(\mu^*) = \begin{cases} 1, & \sqrt{\mu^*} \leq \frac{|b_j|}{\alpha_j + 1}, \\ \frac{|b_j|}{\sqrt{\mu^*}} - \alpha_j, & \frac{|b_j|}{\alpha_j + 1} < \sqrt{\mu^*} < \frac{|b_j|}{\alpha_j}, \\ 0, & \sqrt{\mu^*} \geq \frac{|b_j|}{\alpha_j}. \end{cases} \quad (15)$$

for  $b_j = \|\sqrt{d_j} A_{s_j} t\|_2$  and  $\alpha_j = \lambda d_j (> 0)$ , and  $\tilde{\eta} = \frac{1}{\sqrt{\mu^*}}$  is a root of the function

$$g_t(\eta) \equiv \sum_{i=j}^m u_j(\eta) - k, \quad (16)$$

which is nondecreasing and satisfies

$$g_t \left( \frac{\lambda \cdot \min_{j \in [m]} \{d_j\}}{\|\sqrt{d_1} \|A_{s_1} t\|_2, \sqrt{d_2} \|A_{s_2} t\|_2, \dots, \sqrt{d_m} \|A_{s_m} t\|_2\|_\infty} \right) = \sum_{i=1}^m 0 - k < 0, \quad (17)$$

and,

$$g_t \left( \frac{\lambda \|d_1, d_2, \dots, d_m\|_\infty + 1}{\|M_{\langle s_m \rangle}(\sqrt{d_j} t)\|_2} \right) = \sum_{i=1}^m 1 - k > 0. \quad (18)$$

In addition,  $g_t$  can be reformulated as the sum of pairs of the functions

$$v_i(\eta) \equiv |\eta|b_j| - \alpha_j|, w_i(\eta) \equiv 1 - |\eta|b_j| - (\alpha_j + 1)|, j \in [m], \quad (19)$$

such that,

$$g_t(\eta) = \frac{1}{2} \sum_{j=1}^m v_j(\eta) + \frac{1}{2} \sum_{j=1}^m w_j(\eta) - k. \quad (20)$$

The following important remarks are in order.

**Remark 3** (Root search application for function 16). *Employing the randomized root search method in (Beck & Refael, 2022, Algorithm 1) with the  $2m$  one break point piece-wise linear functions  $v_j, w_j$ , as an input to the algorithm, the root of  $g_t$  can be found in  $O(m)$  time.*

**Remark 4** (Computational complexity of  $\text{prox}_{\lambda \mathcal{GS}_k}$ ). *The computation of  $\text{prox}_{\lambda \mathcal{GS}_k}$  boils down to a root search problem (see, Remark 3), which requires  $O(m)$  operations. In addition, before employing the root search, the assembly of  $v_j, w_j$ , requires the calculations of the  $m$  values of  $b_j$  defined in corollary 16. Note that for any  $j \in [m]$  calculating  $b_j$  is equivalent to  $t^\top A_{s_j} t = \sum_{i \in s_j} t_i^2$ . Since  $s_j$ 's are given, the computational complexity of calculating all  $m$  of  $b_j$  is linear in  $n$ , which is the dimension of  $t$ . Thus, the total computational operations of calculating  $\text{prox}_{\lambda \mathcal{GS}_k}$  summarizes to  $n$  with is the dimension of all groups parameters together.*

## 5 OPTIMIZATION PROCEDURE

The general problem we are solving is of the form

$$\min_{\mathbf{x} \in \mathbb{R}^n} f(\mathbf{x}) + h(\mathbf{x}), \quad (21)$$

<sup>1</sup>If  $A_{s_j} t = 0$ , then equation 15 implies that  $u_i(\mu) = 0$  for all  $\mu \geq 0$ .

**Algorithm 1: General Stochastic Proximal Gradient Method****Input:** Step size  $\alpha_t, \{\rho_t\}_{t=1}^{t=T} \in [0, 1)$ , regularization parameter  $\lambda$ .**Initialization:**  $\theta_1 \in \mathbb{R}^n$  and  $\mathbf{m}_0 = \mathbf{0} \in \mathbb{R}^n$ .**for** iteration  $t = 1, \dots, T$ :    Draw a minibatch sample  $\xi_t$      $\mathbf{g}_t \leftarrow \nabla f(\theta_t; \xi_t)$      $\mathbf{m}_t \leftarrow \rho_t \mathbf{m}_{t-1} + (1 - \rho_t) \mathbf{g}_t$      $\theta_{t+1} \leftarrow \text{prox}_{\alpha_t \lambda h}(\theta_t - \alpha_t \mathbf{m}_t)$ **return**  $\theta$ **Algorithm 2: Learning structured  $k$ -level sparse neural-network by Prox SGD with WGSEF regularization****Input:** Step size  $\alpha_t, \{\rho_t\}_{t=1}^{t=T} \in [0, 1)$ , regularization parameters  $\{\lambda_l\}_{l=1}^{l=L} \in [0, \infty)$ .**Initialization:** Randomly initialize the weights  $\theta_{t=0} \in \mathbb{R}^n, \mathbf{m}_0 = \mathbf{0} \in \mathbb{R}^n$ .**for** iteration  $t = 1, \dots, T$ :    Draw a minibatch sample  $\xi_t$      $\mathbf{g}_t \leftarrow \nabla \mathcal{L}(\theta_t; \xi_t)$      $\mathbf{m}_t \leftarrow \rho_t \mathbf{m}_{t-1} + (1 - \rho_t) \mathbf{g}_t$      $\theta_{t+1} \leftarrow \text{prox}_{\alpha_t \lambda \mathcal{G} \mathcal{S}_k}(\theta_t - \alpha_t \mathbf{m}_t)$ **Prune (Optional):** all  $\#(m - k)$  smallest  $\ell_2$  group norm values of  $\theta$ .**return**  $\theta$ 

where  $f = \frac{1}{N} \sum_{i=1}^N f_i : X \rightarrow \mathbb{R}$  is continuously differentiable, but possibly nonconvex, and  $h$  is a convex function, but possibly nonsmooth. For practical reasons, we cannot store the full gradient  $\nabla f(\mathbf{x})$ . Hence, we would like to use a stochastic gradient type algorithm. However, such a structure poses several difficulties from an optimization perspective, as most research of stochastic first-order algorithms does not account for both nonconvex smooth term and a nonsmooth convex regularizer. In Ghadimi et al. (2016) the authors provide an analysis of a simple stochastic proximal gradient algorithm, where at each iteration a minibatch of weights is updated using a gradient step followed by a proximal step. This algorithm is proved to converge, however, the rate of convergence depends heavily on the minibatch size, and, in fact, for reasonably sized minibatches it will not converge. J Reddi et al. (2016) proposes variance-reduction type algorithms, but since these extend SAGA Defazio et al. (2014) and SVRG Johnson & Zhang (2013) to the nonconvex and nonsmooth setting, they require storing the gradient for each sample (SAGA) which requires  $\mathcal{O}(Nn)$  storage, or recomputing the full gradient every  $s \geq N$  iterations (SVRG), which is undesirable for training neural networks. ProxSVGR+ Li & Li (2018) tackles this issue by calculating the gradient for a batch of  $|B|$  samples. ProxSVGR+ also has a proven convergence rate, however, this rate is strongly dominated by the magnitude of  $|B|$ , hence, impractical for training neural networks.

The ProxSGD algorithm Yang et al. (2020) appears appealing to our problem as it allows for momentum. While the algorithm has a convergence guarantee, the authors do not provide the rate, making it less appealing, given the known issue with the minibatch size. We have found the most suitable optimization algorithm to be ProxGen Yun et al. (2021), as it can accommodate momentum, and also has a proven convergence rate with a fixed and reasonable minibatch size of order  $\Theta(\sqrt{N})$ . Next, we provide a convergence guarantee for Algorithm 1, as given in Yang et al. (2020). This result holds under several regularity assumptions which can be found in Appendix A.3.1.

**Corollary 5.0.1.** *Under Assumptions (C-1)–(C-3) in Appendix A.3.1, Algorithm 1 with constant minibatch size  $b_t = b = \Theta(T)$  is guaranteed to yield  $\mathbb{E}[\text{dist}(\mathbf{0}, \hat{\partial}F(\theta))^2] \leq O(T^{-1})$ , where  $\hat{\partial}F$  is the Fréchet sub-differential function of  $F$  (see, Definition 2 in Appendix A.3.1).*

Next, we propose Algorithm 2, as an implementation of Algorithm 1 to solve equation 8, where  $f \triangleq \mathcal{L}$  and  $h \triangleq \mathcal{G} \mathcal{S}_k$ . Calculating  $\nabla \mathcal{L}(\theta_t; \xi_t)$ , commonly approached using a propagation algorithm, is at least linear in the number of parameters and is obviously getting more complex as the number of layers increases. Therefore, the calculation of the prox of the GSFE is not a bottleneck of the update step complexity (since it is linear in the number of parameters). Notice that the reg-

ularization in our setting is layer-separable as can be seen from the definition of the regularization function  $h(\mathbf{x}) = \sum_{l=1}^L h_l(\mathbf{x}_l)$ . Therefore prox is applied to each layer separately (Beck, 2017, Theorem 6.6). Moreover, we allow different  $\lambda_l, k_l$  parameters per layer. Accordingly, regularization can be applied differently, namely, not in a per-layer fashion. For example, all convolutional layers can be regularized at once, and groups will be defined so that they create cross-layer sparsity. The only condition is that groups are not overlapping. We note that since Assumptions (C-1)–(C-3) are met, Algorithm 2 converges to an  $\epsilon$ -stationary point.

Another technique for solving equation 21 is using the HSPG family of algorithms in Chen et al. (2021; 2020); Dai et al. (2023). These algorithms utilize a two-step procedure in which optimization is carried by standard first-order methods (i.e., subgradient or proximal) to find an approximation that is “sufficiently close” to a solution. This step is followed by a half-space step that freezes the sparse groups and applies a tentative gradient step on the dense groups. Over these dense groups, parameters are zeroed-out if a sufficient decrease condition is met, otherwise, a standard gradient step is executed. Notice that the half-space step, as a variant of a gradient method, requires the regularizer term  $h$  to have a Lipschitz continuous gradient, which is not satisfied in our setting. However, this property is required only for the dense groups as there is no use of the gradient in groups that are already sparse. Since the continuity is violated only for sparse groups, the condition is satisfied in the required region. Finally, while the “sufficiently close” condition mentioned above cannot be verified in practice, simple heuristics to switch between steps still work well. We can either run the first-order step for a fixed number of iterations before switching to the half-space step, or, alternatively, run the first-order step until the sparsity level stabilizes, and then switch to the half-space step. The dense groups are defined as  $\mathcal{I}^0(\mathbf{x}) := \{\gamma \mid \gamma \in \mathcal{G}, \|\mathbf{x}^\gamma\| = 0\}$ , and sparse groups are defined as  $\mathcal{I}^{\neq 0} := \{\gamma \mid \gamma \in \mathcal{G}, \|\mathbf{x}^\gamma\| \neq 0\}$ . The HSPG pseudo-code (proximal-gradient variant) is given in Algorithm 3 in the appendix. We mention the enhanced variant of the standard HSPG algorithm named AdaHSPG+ Dai et al. (2023) improves upon that implementing adaptive strategies that optimize performance, focusing on better handling of complex or dynamic problem scenarios where standard HSPG may be less efficient.

## 6 EXPERIMENTS

We provide only a summary of our experimental results here, deferring full details to Appendix A.5. To demonstrate the performance of prox-SGD with WSEG (Algorithm 2), in terms of both compression and accuracy, we utilized widely-recognized benchmark DNN architectures: VGG16 Simonyan & Zisserman (2014), ResNet18 He et al. (2016a), and MobileNetV1 Howard et al. (2017). These architectures were tested on the datasets CIFAR10 Krizhevsky et al. (2009) and Fashion-MNIST Xiao et al. (2017). All experiments were conducted over 300 epochs. For the first 150 epochs, we employed Algorithm 2, and for the subsequent epochs, we used the HSPG with the WGSEF as a regularize (Algorithm 3). Experiments were conducted using a mini-batch size of  $b = 128$  on an A100 GPU. The coefficient for the WGSE regularizer was set to  $\lambda = 10^{-2}$ . In Table 1, we compare our results with those reported in Dai et al. (2023). The primary metrics of interest are the neural network group sparsity ratio and prediction accuracy (Top-1). Notably, the WGSE achieves a markedly higher group sparsity compared to all other methods except AdaHSPG+, for which we obtained slightly higher results. It should be mentioned that all techniques achieved comparable generalization error rates on the validation datasets. We mention here that all the optimization methods in Table 1 used the well-known Group Lasso as a regularizer, because as observed in Bui et al. (2021), it gave the best results, in terms of compression and performance, in comparison to many other convex and nonconvex regularizers.

Next, we examine the effectiveness of the WGSEF in the LeNet-5 convolutional neural network LeCun et al. (1998), on the MNIST dataset LeCun & Cortes (2010). The networks were trained without any data augmentation. We apply the WGSEF regularization on filters in convolutional layers using a predefined value for the sparsity level  $k$ . Table 2 summarizes the number of remaining filters at convergence, floating-point operations (FLOP), and the speedups. We evaluate these metrics both for a LeNet-5 baseline (i.e., without sparsity learning), and our WGSEF sparsification technique. To be accurate and fair in comparison, the baseline model was trained using SGD. It can be seen that WGSEF reduces the number of filters in the convolution layers by a factor of the half, as dictated by  $k = 8$ , while the accuracy level did not decrease. Furthermore, since the sparsification is structural, there is a significant improvement in flops, as well as in the latency time of inference.



Table 1: Comparison of state-of-the-art techniques to our WGSEF technique, in terms of group sparsity ratio/validation accuracy in percentage, for various models and datasets.

Model	Dataset	Prox-SG	Prox-SVRG	HSPG	AdaHSPG+	WGSEF
VGG16	CIFAR10	54.0/90.6	14.7/89.4	74.6/91.1	76.1/91.0	<b>76.8/91.5</b>
	F-MNIST	19.1/93.0	0.5/92.7	39.7/ <b>93.0</b>	51.2/92.9	<b>51.9/92.8</b>
ResNet18	CIFAR10	26.5/94.1	2.8/94.2	41.6/94.4	42.1/ <b>94.5</b>	<b>42.6/94.5</b>
	F-MNIST	0.0/94.8	0.0/94.6	10.4/ <b>94.9</b>	43.9/ <b>94.9</b>	<b>44.2/94.9</b>
MobileNetV1	CIFAR10	58.1/91.7	29.2/90.7	65.4/ <b>92.0</b>	71.5/91.8	<b>71.8/91.9</b>
	F-MNIST	62.6/94.2	42.0/94.2	74.3/94.5	78.9/ <b>94.6</b>	<b>79.1/94.5</b>

Repeating the same experiment, but now constraining the number of non-pruned filters in the second convolutional layer to be at most 4 (i.e., at most quarter of the baseline), the accuracy slightly deteriorates; however, significant improvements can be observed in both the flops number and the speed up, as expected. The networks were trained with a learning rate of 0.001, regularization magnitude  $\lambda = 10^{-5}$ , and a batch size of 32 for 150 epochs across 5 runs.

Table 2: Results of running algorithm 2, onto redundant filters in LeNet (in the order of conv1-conv2).

LeNet (MNIST)	Error	Filter (sparsity-level)	FLOP	Speedup
Baseline (SGD)	0.84 %	6-16	100 %-100 %	1.00 $\times$ -1.00 $\times$
WGSEF	0.78 %	3-8	48.7 %-21.6 %	2.06 $\times$ -4.53 $\times$
WGSEF	0.89 %	3-4	48.7 %-14.7 %	2.06 $\times$ -7.31 $\times$
LeNet (MNIST)	Error	Parameters	FLOP	Speedup
Unstructured WGSEF	0.76 %	75/(150)-1200/(2400)	68.7 %-59.2 %	1 $\times$ -1 $\times$

We also implemented deep residual networks (ResNet50) He et al. (2016b) and trained it on CIFAR-10 Krizhevsky et al. (2009), applying WGSEF regularization with a predefined sparsity level that is at most equal to half the number of filters at each of the convolutional layers. Again, to be accurate in comparison, the baseline model was trained using SGD, both with an initial learning rate of  $\alpha_0 = 0.01$ , regularization magnitude  $\lambda = 0.3$ , a batch size of 128, and Cosine Annealing learning rate scheduler. The results are similar in nature to those of the MNIST experiment in Table 3.

Table 3: Average results running algorithm 2 on ResNet50 for CIFAR-10 over 5 trials, after completing 110 epochs, using parameters  $\alpha_0 = 0.01$  and  $\lambda = 0.3$ .

Resnet50 (CIFAR-10)	Error	FLOP	Speedup
Baseline (SGD)	7.62 %	100 %	1.00 $\times$
WGSEF	7.4 %	48.7 %	2.06 $\times$

Finally, table 6 presents the results of training both VGG16 and DenseNet40 Huang et al. (2017) on CIFAR100 Krizhevsky et al. (2009), while applying WGSEF regularization with a predefined sparsity level of half the number of channels for VGG16, and 60% of those of the DenseNet40. The baseline model was trained using SGD, both with an initial learning rate of  $\alpha = 0.01$  and regularization magnitude  $\lambda = 0.01$ .

Table 4: Results of running algorithm 2, onto redundant Channels on CIFAR100, over 250 epochs.

DCNN	Model	Error (%)	Pruned Channels	Overall Density
VGG16	Baseline-unpruned	26.28	-	-
	WGSEF	26.46	50%	41.3%
DenseNet40	Baseline-unpruned	25.36	-	-
	WGSEF	25.6	60%	42.8%

## REFERENCES

- On the complementarity of sparse l0 and cel0 regularized loss. *MDPI*, 2021-2023. URL <https://www.mdpi.com>.
- Zeyuan Allen-Zhu, Yuanzhi Li, and Yingyu Liang. Learning and generalization in overparameterized neural networks, going beyond two layers. *Advances in neural information processing systems*, 32, 2019.
- Andreas Argyriou, Rina Foygel, and Nathan Srebro. Sparse prediction with the  $k$ -support norm. *Advances in Neural Information Processing Systems*, 25, 2012.
- Amir Beck. *First-order methods in optimization*. SIAM, 2017.
- Amir Beck and Yehonathan Refael. Sparse regularization via bidualization. *Journal of Global Optimization*, 82(3):463–482, 2022.
- Kevin Bui, Fredrick Park, Shuai Zhang, Yingyong Qi, and Jack Xin. Structured sparsity of convolutional neural networks via nonconvex sparse group regularization. *Frontiers in applied mathematics and statistics*, pp. 62, 2021.
- Tianyi Chen, Guanyi Wang, Tianyu Ding, Bo Ji, Sheng Yi, and Zhihui Zhu. Half-space proximal stochastic gradient method for group-sparsity regularized problem. *arXiv preprint arXiv:2009.12078*, 2020.
- Tianyi Chen, Bo Ji, Tianyu Ding, Biyi Fang, Guanyi Wang, Zhihui Zhu, Luming Liang, Yixin Shi, Sheng Yi, and Xiao Tu. Only train once: A one-shot neural network training and pruning framework. *Advances in Neural Information Processing Systems*, 34:19637–19651, 2021.
- Yu Cheng, Duo Wang, Pan Zhou, and Tao Zhang. A survey of model compression and acceleration for deep neural networks. *arXiv preprint arXiv:1710.09282*, 2017.
- Matthieu Courbariaux, Itay Hubara, Daniel Soudry, Ran El-Yaniv, and Yoshua Bengio. Binarized neural networks: Training deep neural networks with weights and activations constrained to  $\pm 1$  or  $-1$ . *arXiv preprint arXiv:1602.02830*, 2016.
- Yutong Dai, Tianyi Chen, Guanyi Wang, and Daniel Robinson. An adaptive half-space projection method for stochastic optimization problems with group sparse regularization. *Transactions on Machine Learning Research*, 2023.
- Aaron Defazio, Francis Bach, and Simon Lacoste-Julien. Saga: A fast incremental gradient method with support for non-strongly convex composite objectives. *Advances in neural information processing systems*, 27, 2014.
- Lei Deng, Guoqi Li, Song Han, Luping Shi, and Yuan Xie. Model compression and hardware acceleration for neural networks: A comprehensive survey. *Proceedings of the IEEE*, 108(4): 485–532, 2020.
- Emily L Denton, Wojciech Zaremba, Joan Bruna, Yann LeCun, and Rob Fergus. Exploiting linear structure within convolutional networks for efficient evaluation. *Advances in neural information processing systems*, 27, 2014.
- David L Donoho and Michael Elad. Optimally sparse representation in general (nonorthogonal) dictionaries via  $l_1$  minimization. *Proceedings of the National Academy of Sciences*, 100(5):2197–2202, 2003.
- Jianqing Fan and Runze Li. Variable selection via nonconcave penalized likelihood and its oracle properties. *Journal of the American statistical Association*, 96(456):1348–1360, 2001.
- Saeed Ghadimi, Guanghui Lan, and Hongchao Zhang. Mini-batch stochastic approximation methods for nonconvex stochastic composite optimization. *Mathematical Programming*, 155(1):267–305, 2016.
- Yunchao Gong, Liu Liu, Ming Yang, and Lubomir Bourdev. Compressing deep convolutional networks using vector quantization. *arXiv preprint arXiv:1412.6115*, 2014.

- 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 Dally. Learning both weights and connections for efficient neural network. *Advances in neural information processing systems*, 28, 2015b.
- Song Han, Jeff Pool, John Tran, and William J. Dally. Learning both weights and connections for efficient neural networks. *CoRR*, abs/1506.02626, 2015c. URL <http://arxiv.org/abs/1506.02626>.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016a.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778. IEEE, 2016b.
- Andrew G Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*, 2017.
- Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. *arXiv preprint arXiv:1608.06993*, 2017.
- Sashank J Reddi, Suvrit Sra, Barnabas Poczos, and Alexander J Smola. Proximal stochastic methods for nonsmooth nonconvex finite-sum optimization. *Advances in neural information processing systems*, 29, 2016.
- Max Jaderberg, Andrea Vedaldi, and Andrew Zisserman. Speeding up convolutional neural networks with low rank expansions. *arXiv preprint arXiv:1405.3866*, 2014.
- Rie Johnson and Tong Zhang. Accelerating stochastic gradient descent using predictive variance reduction. *Advances in neural information processing systems*, 26, 2013.
- Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- Vadim Lebedev and Victor Lempitsky. Fast convnets using group-wise brain damage, 2015. URL <https://arxiv.org/abs/1506.02515>.
- Vadim Lebedev, Yaroslav Ganin, Maksim Rakhuba, Ivan Oseledets, and Victor Lempitsky. Speeding-up convolutional neural networks using fine-tuned cp-decomposition. *arXiv preprint arXiv:1412.6553*, 2014.
- Yann LeCun and Corinna Cortes. Mnist handwritten digit database, 2010. URL <http://yann.lecun.com/exdb/mnist/>.
- Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. doi: 10.1109/5.726791.
- Jiashi Li, Qi Qi, Jingyu Wang, Ce Ge, Yujian Li, Zhangzhang Yue, and Haifeng Sun. Oicr: Out-in-channel sparsity regularization for compact deep neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 7046–7055, 2019.
- Zhize Li and Jian Li. A simple proximal stochastic gradient method for nonsmooth nonconvex optimization. *Advances in neural information processing systems*, 31, 2018.
- Zhuang Liu, Jianguo Li, Zhiqiang Shen, Gao Huang, Shoumeng Yan, and Changshui Zhang. Learning efficient convolutional networks through network slimming. In *Proceedings of the IEEE international conference on computer vision*, pp. 2736–2744, 2017.
- Yifei Lou, Penghang Yin, Qi He, and Jack Xin. Computing sparse representation in a highly coherent dictionary based on difference of  $\ell_1$  and  $\ell_2$ . *Journal of Scientific Computing*, 64:178–196, 2015.

- Christos Louizos, Max Welling, and Diederik P Kingma. Learning sparse neural networks through  $L_0$  regularization. *arXiv preprint arXiv:1712.01312*, 2017.
- Dmitry Molchanov, Arsenii Ashukha, and Dmitry Vetrov. Variational dropout sparsifies deep neural networks. In *International Conference on Machine Learning*, pp. 2498–2507. PMLR, 2017.
- Balas Kausik Natarajan. Sparse approximate solutions to linear systems. *SIAM journal on computing*, 24(2):227–234, 1995.
- Mohammad Rastegari, Vicente Ordonez, Joseph Redmon, and Ali Farhadi. Xnor-net: Imagenet classification using binary convolutional neural networks. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part IV*, pp. 525–542. Springer, 2016.
- Simone Scardapane, Danilo Comminiello, Amir Hussain, and Aurelio Uncini. Group sparse regularization for deep neural networks. *Neurocomputing*, 241:81–89, 2017.
- Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- Karen Ullrich, Edward Meeds, and Max Welling. Soft weight-sharing for neural network compression. *arXiv preprint arXiv:1702.04008*, 2017.
- J v. Neumann. Zur theorie der gesellschaftsspiele. *Mathematische annalen*, 100(1):295–320, 1928.
- Wei Wen, Chunpeng Wu, Yandan Wang, Yiran Chen, and Hai Li. Learning structured sparsity in deep neural networks, 2016a. URL <https://arxiv.org/abs/1608.03665>.
- Wei Wen, Chunpeng Wu, Yandan Wang, Yiran Chen, and Hai Li. Learning structured sparsity in deep neural networks. *Advances in neural information processing systems*, 29, 2016b.
- Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*, 2017.
- Yang Yang, Yaxiong Yuan, Avraam Chatzimichailidis, Ruud JG van Sloun, Lei Lei, and Symeon Chatzinotas. Proxsgd: Training structured neural networks under regularization and constraints. In *International Conference on Learning Representations (ICLR) 2020*, 2020.
- Jianbo Ye, Xin Lu, Zhe Lin, and James Z Wang. Rethinking the smaller-norm-less-informative assumption in channel pruning of convolution layers. *arXiv preprint arXiv:1802.00124*, 2018.
- Jaehong Yoon and Sung Ju Hwang. Combined group and exclusive sparsity for deep neural networks. In *International Conference on Machine Learning*, pp. 3958–3966. PMLR, 2017.
- Ming Yuan and Yi Lin. Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 68(1):49–67, 2006.
- Jihun Yun, Aurélie C Lozano, and Eunho Yang. Adaptive proximal gradient methods for structured neural networks. *Advances in Neural Information Processing Systems*, 34:24365–24378, 2021.
- Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning (still) requires rethinking generalization. *Communications of the ACM*, 64(3):107–115, 2021.
- Yang Zhou, Rong Jin, and Steven Chu-Hong Hoi. Exclusive lasso for multi-task feature selection. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pp. 988–995. JMLR Workshop and Conference Proceedings, 2010.
- Hui Zou and Trevor Hastie. Regularization and variable selection via the elastic net. *Journal of the royal statistical society: series B (statistical methodology)*, 67(2):301–320, 2005.

## A APPENDIX

### A.1 PROOFS FOR SECTION 3

#### A.1.1 PROOF FOR LEMMA 3.1

*Proof.* Let us define the axillary diagonal positive definite matrix  $D^{n \times n}$ , where the  $D_{i,i}$  entry holds  $D_{i,i} = \sqrt{d_i}, \forall i \in s_j, d_i = d_j$ . Now, consider the following chain of equalities:

$$\begin{aligned}
gs_k^*(\tilde{\theta}) &= \max_{\theta \in \mathbb{R}^n} \{ \langle \tilde{\theta}, \theta \rangle - gs_k(\theta) \} \\
&= \max_{\theta \in \mathbb{R}^n} \left\{ \tilde{\theta}^\top \theta - \frac{1}{2} \sum_{j=1}^m d_j \|\theta_{s_j}\|_2^2 - \delta_{C_k}(\theta) \right\} \\
&= \max_{\substack{\theta \in C_k \\ \forall i \in s_j, d_i = d_j}} \left\{ \tilde{\theta}^\top \theta - \sum_{j=1}^m d_j \cdot \|\theta_{s_j}\|_2^2 \right\} \\
&= \max_{\substack{\theta \in C_k \\ \forall i \in s_j, d_i = d_j}} \left\{ \tilde{\theta}^\top \theta - \frac{1}{2} \theta^\top D^\top D \theta \right\} \\
&\stackrel{(a)}{=} \max_{\substack{t \in \tilde{C}_k \\ \forall i \in s_j, d_i = d_j}} \left\{ \tilde{\theta}^\top D^{-1} t - \frac{1}{2} t^\top t \right\} \\
&= \max_{\substack{t \in \tilde{C}_k \\ \forall i \in s_j, d_i = d_j}} \left\{ -\frac{1}{2} \|t - D^{-1} \tilde{\theta}\|_2^2 + \frac{1}{2} \|D^{-1} \tilde{\theta}\|_2^2 \right\} \\
&= \frac{1}{2} \|D^{-1} \tilde{\theta}\|_2^2 + \max_{\substack{t \in \tilde{C}_k \\ \forall i \in s_j, d_i = d_j}} \left\{ -\frac{1}{2} \|t - D^{-1} \tilde{\theta}\|_2^2 \right\} \\
&= \frac{1}{2} \|D^{-1} \tilde{\theta}\|_2^2 + \max_{\substack{t \in \tilde{C}_k \\ \forall i \in s_j, d_i = d_j}} \left\{ -\frac{1}{2} \sum_{i=1}^n (t_i - (D^{-1})_{ii} \tilde{\theta}_i)^2 \right\} \\
&\stackrel{(b)}{=} \frac{1}{2} \|D^{-1} \tilde{\theta}\|_2^2 + \max_{\substack{t \in \tilde{C}_k \\ \forall i \in s_j, d_i = d_j}} \left\{ -\frac{1}{2} \sum_{j=1}^m \|M_{s_j}(t - D^{-1} \tilde{\theta})\|_2^2 \right\} \\
&= \frac{1}{2} \|D^{-1} \tilde{\theta}\|_2^2 - \frac{1}{2} \sum_{j=m-k}^m \|M_{\langle s_j \rangle}(D^{-1} \tilde{\theta})\|_2^2 \\
&\stackrel{(c)}{=} \frac{1}{2} \sum_{j=1}^k \|M_{\langle s_j \rangle}(D^{-1} \tilde{\theta})\|_2^2 \\
&= \frac{1}{2} \sum_{j=1}^k \frac{1}{d_j} \|M_{\langle s_j \rangle}(\tilde{\theta})\|_2^2
\end{aligned}$$

where, (a) the set  $\tilde{C}_k$  is given by

$$\tilde{C}_k = \{ \theta : \| \| (D\theta)_{s_1} \|^2, \| (D\theta)_{s_2} \|^2 \dots, \| (D\theta)_{s_m} \|^2 \|_0 \leq k \},$$

(b) follows by the fact that the sum of the squares of the coordinates of the input vector in the support of the disjoint subset  $s_1, s_2, \dots, s_m$  that completes the index space  $\bigcup_{i=1}^m s_i = [n]$  is equal to the sum of squares of the coordinates of the original vector, and (c) follows by the fact that  $\sum_{j=m-k}^m \|M_{\langle s_j \rangle}(D^{-1} \tilde{\theta})\|_2^2$  is the sum of square  $\ell_2$ -norm of the  $m - k$  disjoint subsets of  $D^{-1} \tilde{\theta}$  with the smallest  $\ell_2$ -norm, while is the sum of squared  $\ell_2$ -norm of all disjoint subsets of  $D^{-1} \tilde{\theta}$ .  $\square$

## A.1.2 PROOF FOR LEMMA 3.2

*Proof.* We first note that

$$\sum_{i=1}^k \left\| M_{(s_i)}(D^{-1}\tilde{\theta}) \right\|_2 = \max_{u \in B_k} \sum_{i=1}^m u_i \left\| M_{s_i}(D^{-1}\tilde{\theta}) \right\|_2^2, \quad (22)$$

where

$$B_k \triangleq \{u \in \mathbb{R}^m \mid 0 \leq u \leq e, e^\top u \leq k\}. \quad (23)$$

Now, Consider the following chain of inequalities:

$$\begin{aligned} \mathcal{GS}_k(\theta) &= \max_{\tilde{\theta} \in \mathbb{R}^n} \left\{ \langle \theta, \tilde{\theta} \rangle - gs^*(\tilde{\theta}) \right\} \\ &= \max_{\tilde{\theta} \in \mathbb{R}^n} \left\{ \theta^\top \tilde{\theta} - \frac{1}{2} \sum_{i=1}^k \left\| M_{(s_i)}(D^{-1}\tilde{\theta}) \right\|_2^2 \right\} \\ &= \max_{\substack{\tilde{\theta} \in \mathbb{R}^n \\ t = D^{-1}\tilde{\theta}}} \left\{ \theta^\top Dt - \frac{1}{2} \sum_{i=1}^k \left\| M_{(s_i)}(t) \right\|_2^2 \right\} \\ &= \max_{t \in \mathbb{R}^n} \left\{ \theta^\top Dt - \frac{1}{2} \max_{u \in D_k} \sum_{i=1}^m u_i \left\| M_{s_i}(t) \right\|_2^2 \right\} \\ &= \max_{t \in \mathbb{R}^n} \left\{ \theta^\top Dt - \frac{1}{2} \max_{u \in D_k} \left\{ \sum_{j=1}^m u_j \left( t^\top A_{s_j}^\top A_{s_j} t \right) \right\} \right\} \\ &\stackrel{(a)}{=} \max_{t \in \mathbb{R}^n} \left\{ \theta^\top Dt + \frac{1}{2} \min_{\mathbf{u} \in D_k} \left\{ \sum_{j=1}^m (-u_j) \left( t^\top A_{s_j} t \right) \right\} \right\} \\ &= \frac{1}{2} \max_{t \in \mathbb{R}^n} \left\{ \min_{\mathbf{u} \in D_k} \left\{ 2\theta^\top Dt - \sum_{j=1}^m u_j t^\top A_{s_j} t \right\} \right\} \\ &\stackrel{(b)}{=} \frac{1}{2} \min_{\mathbf{u} \in D_k} \left\{ \max_{t \in \mathbb{R}^n} \left\{ 2\theta^\top Dt - \sum_{j=1}^m u_j t^\top A_{s_j} t \right\} \right\} \\ &= \frac{1}{2} \min_{\mathbf{u} \in D_k} \left\{ \sum_{j=1}^m \max_{t \in \mathbb{R}^n} 2\theta^\top A_{s_j} Dt - u_j t A_{s_j} t \right\} \\ &= \frac{1}{2} \min_{\mathbf{u} \in D_k} \left\{ \sum_{j=1}^m \max_{t \in \mathbb{R}^n} 2\sqrt{d_j} \theta^\top A_{s_j} t - u_j t A_{s_j} t \right\} \\ &= \frac{1}{2} \min_{\mathbf{u} \in D_k} \left\{ \sum_{j=1}^m \phi \left( \sqrt{d_j} A_{s_j} \theta, u_j \right) \right\} \\ &= \frac{1}{2} \min_{\mathbf{u} \in D_k} \left\{ \sum_{j=1}^m d_j \phi \left( A_{s_j} \theta, u_j \right) \right\}, \end{aligned} \quad (24)$$

where (a) follows by the fact that  $A_{s_j}$  is a self-adjoint matrix, and (b) follows from the fact that the objective function is concave w.r.t.  $\tilde{y}$  and convex w.r.t  $u$ , and the MinMax Theorem v. Neumann (1928).  $\square$

## A.1.3 PROOF FOR COROLLARY 3.2.1

*Proof.* Directly by expression (34).  $\square$

## A.2 PROOFS FOR SECTION 4

## A.2.1 PROOF OF LEMMA 4.1

*Proof.* Recall that

$$v = \text{prox}_{\lambda \mathcal{G}_{S_k}}(t) = \underset{\theta \in \mathbb{R}^n}{\text{argmin}} \left\{ \lambda \mathcal{G}_{S_k}(\theta) + \frac{1}{2} \|\theta - t\|_2^2 \right\}.$$

Using Lemma 3.2, the above minimization problem can be written as

$$\begin{aligned} & \min_{\mathbf{u} \in D_k} \min_{\theta \in \mathbb{R}^n} \left\{ \Phi(\theta, u, t) \equiv \frac{\lambda}{2} \sum_{j=1}^m d_j \phi(A_{s_j} \theta, u_j) + \frac{1}{2} \|\theta - t\|_2^2 \right\} \\ &= \min_{\mathbf{u} \in D_k} \min_{\theta \in \mathbb{R}^n} \left\{ \Phi(\theta, u, t) \equiv \frac{\lambda}{2} \sum_{j=1}^m d_j \phi(A_{s_j} \theta, u_j) + \frac{1}{2} M_{s_j}(\theta - t)_2^2 \right\} \\ &= \min_{\mathbf{u} \in D_k} \min_{\theta \in \mathbb{R}^n} \left\{ \Phi(\theta, u, t) \equiv \frac{\lambda}{2} \sum_{j=1}^m d_j \phi(A_{s_j} \theta, u_j) + \frac{1}{2} (\theta - t)^\top A_{s_j} (\theta - t) \right\}. \end{aligned} \quad (25)$$

Solving for  $\theta$ , we get that for any  $j \in [m]$ , if  $d_j A_{s_j} \theta \geq 0$  then,

$$\begin{aligned} \frac{d_j \lambda A_{s_j} \hat{\theta}}{u_j} + A_{s_j} (\hat{\theta} - t) &= 0 \\ A_{s_j} \hat{\theta} \left( \frac{d_j \lambda}{u_j} + 1 \right) - A_{s_j} t &= 0 \\ A_{s_j} \left( \hat{\theta} \left( \frac{d_j \lambda}{u_j} + 1 \right) - t \right) &= 0, \end{aligned}$$

meaning that,

$$v_i = \frac{t_i u_j}{\lambda d_j + u_j}, \quad j \in [m], i \in s_j, \quad (26)$$

or, equivalently,

$$A_{s_j} v = \frac{u_j A_{s_j} t}{\lambda d_j + u_j}, \quad j \in [m]. \quad (27)$$

Next, we show that  $u$  is the minimizer of the problem  $\min_{\mathbf{u} \in D_k} \Phi(\theta, u, t)$ . Equation equation 27 also holds when  $u_j = 0$ , since in that case,  $v_i = \hat{\theta}_i = 0$ , for all  $i \in s_j$ . Plugging equation 27 in  $\Phi$ , yields,

$$\begin{aligned} \Phi(\hat{\theta}, u, t) &= \frac{1}{2} \sum_{j=1}^m d_j \left( \lambda \frac{\hat{\theta}^\top A_{s_j} \hat{\theta}}{u_j} \right) + \frac{1}{2} \|\hat{\theta} - t\|_2^2 \\ &= \frac{1}{2} \sum_{j=1}^m \left( \lambda d_j \frac{\hat{\theta}^\top A_{s_j} \hat{\theta}}{u_j} + \|A_{s_j} (\hat{\theta} - t)\|_2^2 \right) \\ &= \frac{1}{2} \sum_{j=1}^m \left( \lambda d_j \frac{u_j^2 t^\top A_{s_j} t}{u_j (\lambda d_j + u_j)^2} + \left\| \frac{\lambda d_j A_{s_j} t}{\lambda d_j + u_j} \right\|_2^2 \right) \\ &= \frac{1}{2} \sum_{j=1}^m \left( \lambda d_j \frac{u_j t^\top A_{s_j} t}{(\lambda d_j + u_j)^2} + \frac{(\lambda d_j)^2 t^\top A_{s_j} t}{(\lambda d_j + u_j)^2} \right) \\ &= \frac{\lambda}{2} \sum_{j=1}^m d_j \frac{t^\top A_{s_j} t}{\lambda d_j + u_j} \\ &= \frac{\lambda}{2} \sum_{j=1}^m \phi \left( \sqrt{d_j} A_{s_j} t, \lambda d_j + u_j \right), \end{aligned} \quad (28)$$

which concludes the proof.  $\square$

## A.2.2 PROOF OF COROLLARY 4.1.1

*Proof.* Assigning a Lagrange multiplier for the inequality constraint  $\mathbf{e}^T \mathbf{u} \leq k$  in problem (25), we obtain the Lagrangian function

$$L(\mathbf{u}, \mu) = \sum_{j=1}^m \left( \phi \left( \sqrt{d_j} A_{s_j} t, \lambda d_j + u_j \right) + \mu u_j \right) - k\mu.$$

Therefore, the dual objective function is given by

$$q(\mu) \equiv \min_{\mathbf{u}: 0 \leq \mathbf{u} \leq \mathbf{e}} L(\mathbf{u}, \mu) = \sum_{j=1}^m \varphi_{b_j, \alpha_j}(\mu) - k\mu, \text{ for, } b_j = \left\| \sqrt{d_j} A_{s_j} t \right\|_2 \text{ and } \alpha_j = \lambda d_j (> 0). \quad (30)$$

where for any  $b \in \mathbb{R}$  and  $\alpha \geq 0$ , the function  $\varphi_{b, \alpha}$  is defined in Beck & Refael (2022) by

$$\varphi_{b, \alpha}(\mu) \equiv \min_{0 \leq u \leq 1} \{ \phi(b, \alpha + u) + \mu u \}, \quad \mu \geq 0. \quad (31)$$

Thus, the dual of problem (14) is the maximization problem

$$\max \{ q(\mu) : \mu \geq 0 \} \quad (32)$$

A direct projection of Lemma (Beck & Refael, 2022, lemma 2.4) is that if  $\tilde{\mu} > 0$ , the function  $\mathbf{u} \mapsto L(\mathbf{u}, \tilde{\mu})$  has a unique minimizer over  $\{\mathbf{u} \in \mathbb{R}^m : \mathbf{0} \leq \mathbf{u} \leq \mathbf{e}\}$  given by  $u_j = \varphi'_{b_j, \alpha_j}(\tilde{\mu})$ , where it was shown that

$$\varphi_{b_j, \alpha_j}(\mu) = \begin{cases} \frac{b_j^2}{\alpha_j + 1} + \mu, & \sqrt{\mu} \leq \frac{|b_j|}{\alpha_j + 1}, \\ 2|b_j|\sqrt{\mu} - \alpha_j\mu, & \frac{|b_j|}{\alpha_j + 1} < \sqrt{\mu} < \frac{|b_j|}{\alpha_j}, \\ \frac{b_j^2}{\alpha_j}, & \sqrt{\mu} \geq \frac{|b_j|}{\alpha_j}, \end{cases}$$

for  $b > 0$ , otherwise 0, and the minimizer is given by

$$u(\mu^*) = \begin{cases} 1, & \sqrt{\mu} \leq \frac{|b_j|}{\alpha_j + 1}, \\ \frac{|b_j|}{\sqrt{\mu}} - \alpha_j, & \frac{|b_j|}{\alpha_j + 1} < \sqrt{\mu} < \frac{|b_j|}{\alpha_j}, \\ 0, & \sqrt{\mu} \geq \frac{|b_j|}{\alpha_j}. \end{cases}$$

Problem (32), is concave differentiable and thus the minimizer  $\tilde{\mu}$  holds  $q'(\tilde{\mu}) = 0$ , meaning

$$q'(\tilde{\mu}) = \sum_{j=1}^m u_j(\mu) - k = 0.$$

We observe that for any  $j \in [m]$  the functions  $u_j(\mu)$  are monotonically continuous nonincreasing, and therefore utilizing Lemma (Beck & Refael, 2022, lemma 3.1)  $\mu^* = \frac{1}{\eta^2}$  is the a root of the nondecreasing function,

$$g_t(\eta) \equiv \sum_{j=1}^m u_j(\eta) - k.$$

Note that for

$$g_t \left( \frac{\lambda \cdot \min_{j \in [m]} \{d_j\}}{\left\| \sqrt{d_1} \|A_{s_1} t\|_2, \sqrt{d_2} \|A_{s_2} t\|_2, \dots, \sqrt{d_m} \|A_{s_m} t\|_2 \right\|_\infty} \right) = \sum_{i=1}^m 0 - k < 0,$$

while

$$g_t \left( \frac{\lambda \|d_1, d_2, \dots, d_m\|_\infty + 1}{\|M_{(s_m)}(\sqrt{d_j} t)\|_2} \right) = \sum_{i=1}^m 1 - k > 0.$$

Now, applying lemma (Beck & Refael, 2022, lemma 3.2), we deduce that  $u_j(\mu)$ , can be divided into the sum of the two following functions,

$$v_j(\eta) \equiv |\eta|b_j - \alpha_j, w_j(\eta) \equiv 1 - |\eta|b_j - (\alpha_j + 1), j \in [m],$$

and thus  $g_t$  can be reformulated as follows

$$g_t(\eta) = \frac{1}{2} \sum_{j=1}^m v_j(\eta) + \frac{1}{2} \sum_{j=1}^m w_j(\eta) - k.$$

□



### A.3 PROOFS FOR SECTION 5

#### A.3.1 ASSUMPTIONS

Consider the following definition needed to present our convergence result.

**Definition 2** (Fréchet Subdifferential). *Let  $\varphi$  be a real-valued function. The Fréchet subdifferential of  $\varphi$  at  $\bar{\theta}$  with  $|\varphi(\bar{\theta})| < \infty$  is defined by*

$$\widehat{\partial}\varphi(\bar{x}) \triangleq \left\{ \theta^* \in \Omega \mid \liminf_{\theta \rightarrow \bar{\theta}} \frac{\varphi(\theta) - \varphi(\bar{\theta}) - \langle \theta^*, \theta - \bar{\theta} \rangle}{\|\theta - \bar{\theta}\|} \geq 0 \right\}. \quad (33)$$

The following are the assumptions under which Corollary 5.0.1 holds.

(C-1) (*L*-smoothness) The loss function  $f$  is differentiable, *L*-smooth, and lower-bounded,

$$\|\nabla f(x) - \nabla f(y)\| \leq L\|x - y\| \quad \text{and} \quad f(x^*) > -\infty.$$

(C-2) (Bounded variance) The stochastic gradient  $g_t = \nabla f(\theta_t; \xi)$  is unbiased and,

$$\mathbb{E}_\xi[\nabla f(\theta_t; \xi)] = \nabla f(\theta_t), \quad \mathbb{E}_\xi[\|g_t - \nabla f(\theta_t)\|^2] \leq \sigma^2.$$

(C-3) (i) Final step-vector is finite, (ii) the stochastic gradient is bounded, and (iii) the momentum parameter is exponentially decaying, namely,

$$(i) \|\theta_{t+1} - \theta_t\| \leq D, \quad (ii) \|g_t\| \leq G, \quad (iii) \rho_t = \rho_0 \mu^{t-1},$$

with  $D, G > 0$  and  $\rho_0, \mu \in [0, 1)$ .

#### A.3.2 PROOF OF COROLLARY 5.0.1

*Proof.* The proof follows from (Yun et al., 2021, Corollary 1), by taking  $C_t = \mathbf{0}$  and  $\delta = 1$ .  $\square$

### A.4 ADDITIONAL ALGORITHM

---

#### Algorithm 3: General Stochastic Proximal Gradient Method

---

**Input:** Stepsize  $\alpha_t$ ,  $\{\rho_t\}_{t=1}^T \in [0, 1)$ , regularization parameter  $\lambda$ , switch condition  $\mathcal{S}$ , projection threshold  $\epsilon$ .

**Initialization:**  $\theta_1 \in \mathbb{R}^n$  and  $\mathbf{m}_0 = \mathbf{0} \in \mathbb{R}^n$ .

**for** iteration  $t = 1, \dots, T$ :

**if** condition  $\mathcal{S}$  is **not** satisfied:

    Apply Algorithm 2

**else:**

    Draw a minibatch sample  $\xi_t$

$\mathbf{g}_t \leftarrow \nabla f(\theta_t^{\mathcal{I}^{\neq 0}}; \xi_t) + \nabla h(\theta_t^{\mathcal{I}^{\neq 0}}; \xi_t)$

$\tilde{\theta}_t^{\mathcal{I}^{\neq 0}} \leftarrow \theta_{t-1} - \alpha_t \mathbf{g}_t$ ,  $\tilde{\theta}_t^{\mathcal{I}^0} \leftarrow \mathbf{0}$

**for** each group  $\gamma \in \mathcal{I}^{\neq 0}$ :

**if**  $\langle \tilde{\theta}_t^\gamma, \theta_{t-1}^\gamma \rangle < \epsilon \|\theta_{t-1}^\gamma\|^2$ :

$\tilde{\theta}_t^\gamma \leftarrow \mathbf{0}$

$\theta_{t+1} \leftarrow \tilde{\theta}_t$

**return**  $\theta$

---

### A.5 ADDITIONAL EXPERIMENTS

We examine the effectiveness of the WGSEF in the LeNet convolutional neural network LeCun et al. (1998) (the architecture is given in the appendices), on the FashionMNIST dataset. The networks were trained without any data augmentation. We apply the WGSEF regularization on filters in convolutional layers using a predefined value for the sparsity level  $k$ . Table 5 summarizes the number of remaining filters at convergence, floating-point operations (FLOP), and the speedups. We evaluate these metrics both for a LeNet baseline (i.e., without sparsity learning), and our WGSEF sparsification technique. To be accurate and fair in comparison, the baseline model was trained using SGD. We use a learning rate equal to  $1e-4$ , with a batch size of 32, a momentum 0.95, and 15 epochs.


Table 5: Results of training while applying WGSEF sparsification (with  $\lambda = 0.05$ ), onto redundant filters in LeNet (in the order of conv1-conv2), and neurons in Linear layers.

LeNet (FasionMNIST)	Error	Filter (non-sparse)	Linear layers sparsity	Speedup (conv1-conv2)
Baseline (SGD)	11.1 %	6-16	0 %	1.00 $\times$ -1.00 $\times$
WGSEF	11 %	3-8	0%	2 $\times$ -4.5 $\times$
WGSEF	14 %	4-6	62%	1.7 $\times$ -6.1 $\times$
WGSEF	12.3 %	2-3	0%	2 $\times$ -7.12 $\times$

The second row of Table 5 shows that the method resulted in a significant decrease in the number of non-zero filters, in this case, half of the filters (groups of parameters) were zeroed while still improving the model performance. The rest of the experiments show that there was a higher sparsification in the number of non-zero filters (groups of parameters), with only a negligible degradation in the model’s accuracy.

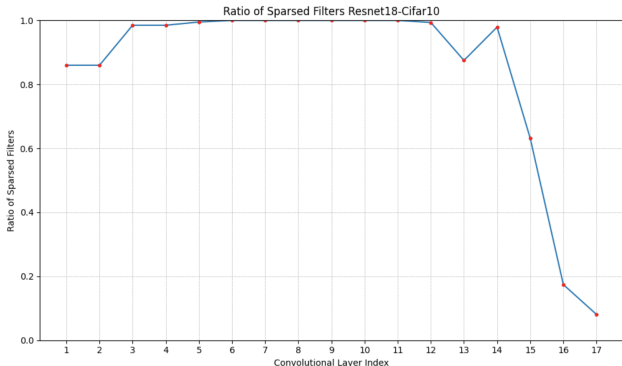
Following, in figure 1, the two graphs refer to the training procedure of the model corresponding with the last row of Table 5. These graphs illustrate the sparsity of the model parameters (in percentages) and the corresponding accuracy as a function of the epoch number. Notably, the desired sparsity level was rapidly attained within the first three epochs. Subsequently, the model enhanced its accuracy throughout the remaining epochs without compromising the achieved sparsity.

Figure 1: The graph on the left shows the level of sparseness (in percentages) as a function of epoch number, while the graph on the right shows the model’s accuracy as a function of epoch number.



This shows that Algorithm 2 results in the actual zeroing of the (groups of) parameters intended for pruning, and not just pushing them close to zero, hence practically the last line in Algorithm 2 which proposed as optional, is actually barely needed.

Figure 2: Ratio of the sparse filters in the convolutional layers, according to the order of the layers in the resnet18 model, as obtained by Algorithm 2, and corresponds to the experiment in row 3 of table 2.



## A.6 LEnet CONVOLUTIONAL NEURAL NETWORK ARCHITECTURE

Layer	# filters / neurons	Filter size	Stride	Size of feature map	Activation function
Input	—	—	—	$32 \times 32 \times 1$	
Conv 1	6	$5 \times 5$	1	$28 \times 28 \times 6$	Relu
MaxPool2d		$2 \times 2$	2	$14 \times 14 \times 6$	
Conv 2	16	$5 \times 5$	1	$10 \times 10 \times 16$	Relu
MaxPool2d		$2 \times 2$	2	$5 \times 5 \times 16$	
Fully Connected 1	—	—	—	120	Relu
Fully Connected 2	—	—	—	84	Relu
Fully Connected 3	—	—	—	10	Softmax