

# COPS: Controlled Pruning Before Training Starts

Wimmer Paul

Image Processing

Robert Bosch GmbH & Lübeck University

71229 Leonberg, Germany

{Paul.Wimmer,

Mehnert Jens

Image Processing

Robert Bosch GmbH

71229 Leonberg, Germany

JensEricMarkus.Mehnert,

Condurache Alexandru

Engineering Cognitive Systems

Robert Bosch GmbH & Lübeck University

70499 Stuttgart, Germany

AlexandruPaul.Condurache}@de.bosch.com

**Abstract**—State-of-the-art deep neural network (DNN) pruning techniques, applied one-shot before training starts, evaluate sparse architectures with the help of a single criterion—called pruning score. Pruning weights based on a solitary score works well for some architectures and pruning rates but may also fail for other ones. As a common baseline for pruning scores, we introduce the notion of a generalized synaptic score (GSS). In this work we do not concentrate on a single pruning criterion, but provide a framework for combining arbitrary GSSs to create more powerful pruning strategies. These *CO*mbined *PR*uning *SC*ores (COPS) are obtained by solving a constrained optimization problem. Optimizing for more than one score prevents the sparse network to overly specialize on an individual task, thus *CO*ntrols *PR*uning before training *St*arts. The combinatorial optimization problem given by COPS is relaxed on a linear program (LP). This LP is solved analytically and determines a solution for COPS. Furthermore, an algorithm to compute it for two scores numerically is proposed and evaluated. Solving COPS in such a way has lower complexity than the best general LP solver. In our experiments we compared pruning with COPS against state-of-the-art methods for different network architectures and image classification tasks and obtained improved results.

## I. INTRODUCTION

*Network pruning* [1]–[3] sets parts of a DNN’s weights to zero. This can help to reduce the model’s complexity and memory requirements, speed up inference [4] and even lead to better generalization ability for the network [2]. In recent years, *training* sparse networks, *i.e.* networks with many weights fixed at zero, became of interest to the deep learning community [5]–[9], providing the potential benefits of reduced runtime and memory requirements not only for inference but also for training. Sparse networks induce sparse gradient computations, but also the problem of weak gradient signals if too many parameters are pruned [8], [9]. One possibility to maintain a strong gradient signal while training only sparse parts of the network can be achieved by *freezing* big portions of the weights during training at their initial value [10]. But likewise, pruned networks can have a sufficient gradient flow, even for high pruning rates, if the pruning criterion is chosen cautiously [8], [9]. In this work we will focus on one-shot pruning methods, applied before training starts. By one-shot pruning we mean pruning in a single step, not iteratively. Thus, no resources are spent for pre-training [6], iterative pruning [8] or a dynamical change of the network during training [5]. The two state-of-the-art (SOTA) pruning methods applied one-




Fig. 1. Comparison of dense baseline, random pruning, SNIP, GraSP and G-COPS for a VGG16 on CIFAR-10. Left: Mean gradient norm of the network’s remaining weights, evaluated before training and normalized w.r.t. the dense network’s gradient. Right: Corresponding test accuracies, see also Table III.

shot before training are *Single-Shot Network Pruning based on Connection Sensitivity* (SNIP) [7] and *Gradient Signal Preservation* (GraSP) [9]. For some conditions SNIP achieves better results than GraSP, for others GraSP outperforms SNIP as shown for a VGG16 [11] trained on CIFAR-10 [12] in Fig. 1, right-hand side. SNIP trains those weights having the biggest individual influence in changing the loss function at the beginning of training [7], the weights with the highest *saliency*. However, choosing weights only based on their solitary high saliency does not guarantee a sufficient information flow in the sparse network [9]. For high pruning rates, this usually leads to a reduced gradient flow, visible in Fig. 1 for the pruning rate  $p = 0.995$ , and finally to the pruning of whole layers [8]. Here and in the following, gradient flow denotes the strength of the gradient signal. Overcoming the low gradient flow for small pruning rates was a motivation for the GraSP method. GraSP zeroes those weights having the smallest *importance*, an approximation of the impact of a weight’s removal on the sparse network’s gradient flow. Setting a weight with positive importance to zero most likely decreases the gradient flow. Thus, GraSP pruned networks have a sufficiently strong gradient signal also for high pruning rates, as shown in Fig. 1. But for lower pruning rates, where SNIP’s gradient flow is strong enough, SNIP leads to better results than GraSP as it explicitly models the effect of pruning on the networks ability to optimize the loss function. Summarized, a higher gradient flow does not necessarily induce a better performance. On the other hand, a gradient flow larger than zero is a necessary condition to train sparse networks successfully. Combining

TABLE I  
MATHEMATICAL NOTATION USED IN THIS WORK.

|   |   |
|---|---|
| $i$ -th Unit Vector   | $e^{(i)} := (\delta_{i,j})_{j=1}^d \in \mathbb{R}^d$ with $\delta_{i,j} = 1$ if $i = j$ , else $\delta_{i,j} = 0$             |
| $i$ -th Vector Component  | $x_i := \langle x, e^{(i)} \rangle := x^T \cdot e^{(i)}$ for $x \in \mathbb{R}^d, i \in \{1, \dots, d\}$                      |
| Vector Support  | $\text{supp } x := \{i \in \{1, \dots, d\} : x_i \neq 0\}$ for $x \in \mathbb{R}^d$   |
| $q$ -Norm   | $\ x\ _0 := \#\text{supp } x$ and $\ \cdot\ _q$ $q$ -norm for $q \in (0, \infty]$   |
| Pruning rate $p$ and Sparsity $\sigma$  | $p := 1 - d^{-1} \cdot \ x\ _0$ and $\sigma = d \cdot (1 - p) = \ x\ _0$ for $x \in \mathbb{R}^d$                             |
| Interior $A^\circ$ of $A \subset \mathbb{R}^d$  | $A^\circ := \{x \in A : \exists \varepsilon > 0 \text{ with } \{y \in \mathbb{R}^d : \ y - x\ _2 < \varepsilon\} \subset A\}$ |
| $n - \text{idxmin}\{\alpha_i : i \in I\}$ for $I$ finite index set, $\alpha_i \in \mathbb{R}, n \in \mathbb{N}$ | $\{J \subset I : \#J = \min\{\#I, n\} \text{ and } \alpha_j \leq \alpha_i \forall j \in J, i \in I \setminus J\}$             |

both, SNIP’s focus on the weights’ saliencies together with a strong information flow, guaranteed by GraSP, motivated us to merge them to achieve better results for both, high and low pruning rates. In this work, all COPS computations involve a SNIP score and a second pruning score that might change between experiments. Thus, the COPS combination of SNIP and GraSP is called G-COPS. Fig. 1 compares pruning with G-COPS against SNIP and GraSP. For all sparsities, G-COPS performs better than both SOTA methods SNIP and GraSP. Particularly for high pruning rates, training weights with high saliencies while guaranteeing a strong gradient flow via G-COPS improves GraSP and SNIP considerably.

#### A. Main contributions

Our main contributions presented in this work are:

- Combining arbitrary generalized synaptic score based pruning methods via COPS.
- Solving the resulting constrained, combinatorial optimization problem analytically by relaxing it on a LP.
- Providing an algorithm for solving the relaxed LP with lower complexity than the best known general LP solver [13].
- Generating one-shot better performing sparse architectures for training than the two SOTA methods SNIP and GraSP by combining and balancing them properly.
- Thereby, COPS is shown to be more efficient than naively combining two pruning scores linearly.

Figures in this work are best viewed in the colored online version. Table I summarizes the mathematical notation.

## II. RELATED WORK

*a) Model compression:* can be achieved by methods such as *quantization, weight sharing, tensor decomposition, low rank tensor approximation* or *pruning*. Quantization reduces the number of bits used to represent the network’s weights and/or activation maps [1]. 32bit floats are replaced by low precision integers, thus decreasing memory consumption and speeding up inference. Memory reduction and speed up can also be achieved by weight sharing [14], tensor decomposition [15] or low rank tensor approximation [16] to name only a few.

*b) Pruning:* is generally distinguished between *structured* and *unstructured* pruning [4]. Structured pruning deletes whole channels, neurons or even coarser structures, immediately resulting in reduced computation time. Unstructured pruning zeroes weights individually. Therefore, better results can be achieved and pruning to higher sparsity levels is

possible, compared to structured pruning [17]. Setting single weights to zero does not automatically lead to a decreased number of computations. Specialized soft- and hardware [18] is needed to obtain also benefits in computational time. In this work, we evaluate COPS only on unstructured pruning methods as they often serve as foundations for corresponding structured methods, see for example [17]. But the theory derived in this paper also works for structured pruning, based on score functions, without the need of any further modification.

Pruned architectures can be created for instance by penalizing non-zero weights during training [19], magnitude pruning [1], [6] or saliency based pruning. For the latter, the significance of weights is measured with the Hessian of the loss [2] or the sensitivity of the loss with respect to inclusion/exclusion of each weight [3].

*c) Training sparse networks:* successfully from scratch was demonstrated by the Lottery Ticket Hypothesis [6]. The so trained networks, called *winning tickets*, can reach the same performance as the baseline architecture up to a high sparsity regime. But to find these winning tickets, many iterative pre-training and pruning steps are needed [6]. Well trainable sparse networks can also be found without costly pre-training via ranking saliencies of weights [7], [20] or preserving the dense network’s information flow for the sparse architecture [8], [9]. This is done either one-shot [7], [9], [20] or iteratively [8], [21]. *Dynamic sparse training* [5] trains sparse networks but enables the sparse architectures to change during training. In this work, we focus on pruning methods applied one-shot before training starts. But for other scoring based pruning methods needing iterative pruning steps [6], [8], [21], or which are applied later on in training [1]–[3], our method can also be used without modifications.

*d) Linear programming:* In order to combine pruning scores, a combinatorial, constrained optimization problem is solved by relaxing the  $\|\cdot\|_0$ -“norm” to the  $\|\cdot\|_1$ -norm. The resulting relaxed problem is shown to be a LP. The dual problem of the relaxed problem is solved analytically with help of convex optimization methods [22]. In practice, contrarily to standard Simplex methods [23], we obtain the solution not by walking between vertices of the polytope, but by using the simple nevertheless robust bisection of intervals. Thus, no pivoting rules are needed to overcome worst case scenarios. LPs can also be solved fast and robustly with *interior point methods* [13].

### III. COMBINED PRUNING SCORE

In this section we introduce the COPS mask. The COPS mask is defined as a solution of a constrained optimization problem which optimizes the *target score* function  $S_0$  over *pruning masks*  $m \in \{0, 1\}^D$ , respecting a sparsity constraint  $\|m\|_0 \leq \sigma$ , while being controlled by a constraint  $\kappa$  on the *control score* function  $S_1$ .

To be consistent with standard notation in convex optimization literature [22], pruning scores are *minimized* in this work. For instance, a high saliency/importance corresponds to a low SNIP/GraSP score, respectively. This is achieved by taking the negative of the original SNIP/GraSP score.

#### A. Basic assumptions and problem formulation

Let  $f_\Theta$  be a DNN with vectorized weights  $\Theta \in \mathbb{R}^D$ . Pruning can be modelled by superimposing a pruning mask  $m \in \{0, 1\}^D$  over the weights via  $m \odot \Theta = (m_i \cdot \Theta_i)_{i=1}^D$ . Here,  $\odot$  denotes the *Hadamard product*. If a component  $m_i$  of the pruning mask is equal to zero, the corresponding weight  $\Theta_i$  will be pruned. If  $m_i = 1$ , the weight  $\Theta_i$  will be active.

SOTA pruning methods applied without any pre-training use, up to changed signs,  $S(m) = -\langle \frac{\partial L}{\partial \Theta} \odot \Theta, m \rangle$  (SNIP), or  $S(m) = -\langle \left( \frac{\partial^2 L}{\partial \Theta^2} \cdot \frac{\partial L}{\partial \Theta} \right) \odot \Theta, m \rangle$  (GraSP) as scores. More general in [8], so called *synaptic scores*  $S(m) = \langle \frac{\partial R}{\partial \Theta} \odot \Theta, m \rangle$  are introduced. Here,  $R$  is a function depending on the weights  $\Theta$  which does not need to be the network's loss function  $L$ . The derivatives of  $L$  and  $R$  are approximated, if necessary, by a sufficient number of training data [7]–[9]. For all three methods, the *score*  $S(m)$  should indicate the performance of the pruned network  $f_{m \odot \Theta}$ , concerning a given criterion. Performance criteria are for example the network's gradient flow for GraSP or the ability to change the loss function for SNIP. All scores above are obtained by evaluating a linear score function  $S$  on a pruning mask  $m$ . We call such scores *generalized synaptic scores* (GSS) and  $S$  a *GSS function*. In the following, we are only interested in the evaluation  $m \mapsto S(m)$ . Thus, we assume the score function  $S$  to be known and ignore potential dependencies, such as from  $\Theta$  or  $L$ , in the notation. For a GSS it holds  $S(m) = \sum_{i=1}^D m_i \cdot S(e^{(i)})$  by linearity. If  $\Theta_i$  is active, the score  $S(e^{(i)})$  can be seen as the contribution of weight  $\Theta_i$  to the overall score of the network. As  $S(m)$  is minimized in the following, pruning a weight  $\Theta_i$  with high contribution  $S(e^{(i)})$  is assumed to lead to better results than pruning a weight with small contribution. Therefore, the goal for pruning to sparsity  $\sigma \in \{1, \dots, D\}$  with a single score function  $S_0$  is given by the optimization problem

$$\min_{m \in \mathcal{X}_{0,\sigma}} S_0(m), \quad (1)$$

with

$$\mathcal{X}_{0,\sigma} := \{m \in \{0, 1\}^D : \|m\|_0 \leq \sigma\}. \quad (2)$$

For a GSS function  $S_0$ , (1) is solved by  $m^* \in \{0, 1\}^D$  with  $\text{supp } m^* \in \sigma - \text{idxmin}\{i \in \{1, \dots, D\} : S_0(e^{(i)}) < 0\}$ . (3)

As shown in Fig. 1, optimizing a single score does not provide the best results for all situations. By knowing the weakness of a score function, as the potential of a small gradient flow for SNIP, we can control it by constraining it. This is modelled by a constraint  $\kappa \in \mathbb{R}$  on the control score function  $S_1$ . The resulting problem for GSS functions  $S_0$  and  $S_1$  is given by

$$\min_{m \in \mathcal{X}_{0,\sigma}} S_0(m), \text{ such that } S_1(m) \leq \kappa. \quad (\mathcal{P}_0)$$

In this form, since  $\mathcal{X}_{0,\sigma}$  is a discrete set,  $(\mathcal{P}_0)$  is a combinatorial problem and cannot be solved with convex optimization methods. Relaxing  $\mathcal{X}_{0,\sigma}$  on the convex optimization set

$$\mathcal{X}_{1,\sigma} = \{m \in [0, 1]^D : \|m\|_1 \leq \sigma\} \supset \mathcal{X}_{0,\sigma} \quad (4)$$

leads to the relaxed, linear problem

$$\min_{m \in \mathcal{X}_{1,\sigma}} S_0(m), \text{ such that } S_1(m) \leq \kappa. \quad (\mathcal{P}_1)$$

#### B. Solution to $(\mathcal{P}_1)$ solves $(\mathcal{P}_0)$

As  $\mathcal{X}_{1,\sigma} = \{m \in [0, 1]^D : \|m\|_1 \leq \sigma\}$  and the constraint on  $S_1$  are given by linear inequalities, we could solve  $(\mathcal{P}_1)$  directly with a LP solver [13], [23]. But in order to show that  $(\mathcal{P}_1)$  yields a solution to  $(\mathcal{P}_0)$ , we will solve the corresponding dual problem. The Lagrangian of  $(\mathcal{P}_1)$  is given by [22]

$$\Lambda(m, \lambda) := S_0(m) + \lambda \cdot (S_1(m) - \kappa) \quad (5)$$

for  $\lambda \in \mathbb{R}$ . With the help of the Lagrangian (5), the dual function of problem  $(\mathcal{P}_1)$  can be defined as

$$g(\lambda) := \inf_{m \in \mathcal{X}_{1,\sigma}} \Lambda(m, \lambda). \quad (6)$$

With  $m \in \mathcal{X}_{1,\sigma}$  and the linearity of  $S_0$  and  $S_1$ ,

$$g(\lambda) = \sum_{i \in I_\lambda} s_i^{(0)} + \lambda \cdot s_i^{(1)} - \lambda \cdot \kappa, \quad s_i^{(k)} := S_k(e^{(i)}), \quad (7)$$

holds for every  $\lambda \geq 0$  and arbitrary

$$I_\lambda \in \sigma - \text{idxmin}\{i \in \{1, \dots, D\} : s_i^{(0)} + \lambda \cdot s_i^{(1)} < 0\}. \quad (8)$$

The dual function  $g$  is needed to solve the dual problem

$$\sup_{\lambda \geq 0} g(\lambda). \quad (\mathcal{D}_1)$$

**Theorem 1.** *Let  $\sigma \in \{1, \dots, D\}$ ,  $S_0$  and  $S_1$  be GSS. Further, assume  $\kappa$  is chosen such that Slater's condition holds, meaning a  $m \in \mathcal{X}_{1,\sigma}^o$  with  $S_1(m) < \kappa$  exists. Then the dual problem  $(\mathcal{D}_1)$  w.r.t.  $(\mathcal{P}_1)$  is solved by a  $\lambda^* \geq 0$ . Let  $\lambda^*$  be an optimal point of the dual function  $g$  and  $I_{\lambda^*}$  corresponding indices of active scores,  $m^*$  computed according to (8). Then*

$$m^* \in \{0, 1\}^D \text{ with } \text{supp } m^* = I_{\lambda^*} \quad (9)$$

*solves the primal problem  $(\mathcal{P}_1)$ . Finally, it holds  $m^* \in \mathcal{X}_{0,\sigma} \subset \mathcal{X}_{1,\sigma}$  and therefore  $m^*$  is also a solution to the un-relaxed problem  $(\mathcal{P}_0)$ .*

*Proof.* By (7) and (8),  $g$  is continuous. For  $\lambda$  big enough, Slater's condition implies that  $g$  is strictly monotonically decreasing. This implicates the existence of a solution  $\lambda^* \geq 0$




Fig. 2. Similarity between masks derived by G-COPS and the underlying target score SNIP ( $\text{sim}(M_0, m_{\alpha_\kappa})$ ) and control score GraSP ( $\text{sim}(M_1, m_{\alpha_\kappa})$ ), respectively. SNIP and GraSP are compared by  $\text{sim}(M_0, M_1)$ .

for  $(\mathcal{D}_1)$ . Strong duality induces complementary slackness (a) [22]. Therefore, for  $m^*$  defined as in (9),  $m^* \in \mathcal{X}_{1,\sigma}$  and

$$g(\lambda^*) \stackrel{(a)}{=} \sum_{i \in I_{\lambda^*}} s_i^{(0)} = S_0(m^*) \quad (10)$$

is true. Meaning that  $m^*$  solves  $(\mathcal{P}_1)$  by strong duality [22]. Finally,  $\mathcal{X}_{0,\sigma} \subset \mathcal{X}_{1,\sigma}$  and  $m^* \in \mathcal{X}_{0,\sigma}$  thus solves  $(\mathcal{P}_0)$ .  $\square$

### C. How to choose the constraint $\kappa$

Before we present an algorithm solving the dual problem  $(\mathcal{D}_1)$ , thus by Theorem 1 also  $(\mathcal{P}_0)$ , numerically, we need to discuss the constraint parameter  $\kappa$ . From now on we assume  $\kappa < 0$ , as  $\kappa$  is chosen negative in all experiments considered in Section IV. Theorem 1 suggests using  $\kappa$  such that Slater's condition holds. With  $s_i^{(k)} = S_k(e^{(i)})$  and the linearity of  $S_k$ ,  $k = 0, 1$ , the constraint  $\kappa$  consequently should be chosen as

$$\kappa > \kappa_{\min} := \sum_{i \in I} s_i^{(1)}, \quad I \in \sigma - \text{idxmin}\{i : s_i^{(0)} < 0\}. \quad (11)$$

Obviously, there is no choice of a pruned network with at most  $\sigma$  remaining weights having a total  $S_1$  score lower than  $\kappa_{\min}$ . If  $\kappa$  is chosen close to  $\kappa_{\min}$ , the solution to  $(\mathcal{P}_0)$  will be restricted strictly towards fulfilling the constraint. Therefore, the resulting network will be close to the one obtained by only taking the control scores  $s_1^{(1)}, \dots, s_D^{(1)}$  into account. Using a constraint  $\kappa \gg \kappa_{\min}$  on the other hand leads to less control. Accordingly, minimizing the target score function  $S_0$  becomes more important. Choosing  $\kappa \geq \kappa_{\max}$  with

$$\kappa_{\max} := \max_{I \in \sigma - \text{idxmin}\{i : s_i^{(0)} < 0\}} \sum_{i \in I} s_i^{(1)} \quad (12)$$

leads to the same solution as the unconstrained problem (1), considering only scores  $s_1^{(0)}, \dots, s_D^{(0)}$ .

Fig. 2 shows the similarity between the SNIP mask  $M_0$ , the GraSP mask  $M_1$  and the combined G-COPS masks  $m_{\alpha_\kappa}$  calculated for varying values for the constraint  $\kappa := \alpha_\kappa \cdot \kappa_{\min}$  for a VGG16 on CIFAR-10. The G-COPS masks are obtained by using SNIP as target and GraSP as control score. As  $\kappa_{\min} < 0$  holds,  $\kappa_{\min} < \alpha_\kappa \cdot \kappa_{\min}$  is true for all  $\alpha_\kappa \in [0, 1)$ . The similarity  $\text{sim}(m, M)$  between two pruning masks  $m, M \in \{0, 1\}^D \setminus \{0\}$  is defined as

$$\text{sim}(m, M) := \frac{\|m \odot M\|_0}{\max\{\|m\|_0, \|M\|_0\}} \in [0, 1]. \quad (13)$$




Fig. 3. Left: Graphical example of  $g$  and  $\lambda^*$  with sparsity  $\sigma = 2$  and underlying  $\pi_i(\lambda) = s_i^{(0)} + \lambda \cdot s_i^{(1)}$ . For clearness,  $\kappa = 0$  is used. Active  $\pi_i(\lambda)$  are colored red. Gray coloring of  $\pi_i(\lambda)$  for a given  $\lambda$  highlights  $\pi_i(\lambda) \geq 0$ . Right: Shows first three steps of Algorithm 1 to obtain a solution  $\lambda^*$ .

It is zero if and only if  $\text{supp } m \cap \text{supp } M = \emptyset$  and one only for  $\text{supp } m = \text{supp } M$ . As suspected, Fig. 2 shows that  $m_{\alpha_\kappa} \rightarrow M_1$  for  $\alpha_\kappa \rightarrow 1$ . Also, for small values of  $\alpha_\kappa$  it holds  $m_{\alpha_\kappa} \rightarrow M_0$ . In between, an increasing of the similarity  $\text{sim}(M_{1-i}, m_{\alpha_\kappa})$  implies a decreasing of  $\text{sim}(M_i, m_{\alpha_\kappa})$ ,  $i = 0, 1$ . Therefore,  $m_{\alpha_\kappa}$  can be seen as an interpolation between the masks  $M_0$  and  $M_1$ , as desired. We additionally observe  $\text{sim}(M_0, M_1) > 0$ , i.e. some indices are chosen for both scores. By construction of the G-COPS mask, these similar weights are also chosen for  $m_{\alpha_\kappa}$  for all  $\alpha_\kappa < 1$ . The similarity between SNIP and GraSP decreases with an increasing pruning rate (gray line in Fig. 2 left plot  $p = 0.9$  compared to right plot  $p = 0.995$ ). Thus, the most important weights for SNIP are less important for GraSP and vice-versa, explaining SNIP's loss of gradient flow for high pruning rates.

### D. Algorithm to compute COPS mask

#### Algorithm 1 COPS

**Require:** Sparsity constraint  $\sigma \in \{1, \dots, D\}$ ,  $\alpha_\kappa \in [0, 1)$ , target scores  $s_1^{(0)}, \dots, s_D^{(0)}$ , control scores  $s_1^{(1)}, \dots, s_D^{(1)}$ , precision  $\delta > 0$ , upper bound  $\lambda_+ > 0$  with  $g'(\lambda_+) \leq 0$

- 1: Set constraint  $\kappa = \alpha_\kappa \cdot \kappa_{\min}$   $\triangleright \kappa_{\min}$  according to (11)
- 2: Calculate  $I_0$  as in (8)  $\triangleright I_0$  gives  $g(0)$  and  $g'(0)$  via (7), (14)
- 3: **if**  $g'(0) \leq 0$  **then**  $\triangleright g$  is monotonically decreasing for  $\lambda \geq 0$
- 4:  $\lambda^* \leftarrow 0$
- 5: **else**  $\triangleright \lambda^* \geq 0, \lambda_- = 0$  is lower bound
- 6: Calculate  $I_{\lambda_+}$   $\triangleright I_{\lambda_+}$  gives  $g(\lambda_+)$  and  $g'(\lambda_+)$
- 7:  $\lambda_-^{(0)} \leftarrow 0, \lambda_+^{(0)} \leftarrow \lambda_+, k \leftarrow 0$   $\triangleright k$  counts steps
- 8: **repeat**  $\triangleright$  Bisection
- 9:  $\lambda_0 \leftarrow \frac{\lambda_-^{(k)} + \lambda_+^{(k)}}{2}$   $\triangleright$  Midpoint  $\lambda_0$
- 10: Calculate  $I_{\lambda_0}$   $\triangleright I_{\lambda_0}$  gives  $g(\lambda_0)$  and  $g'(\lambda_0)$
- 11: **if**  $g'(\lambda_0) = 0$  **then**  $\triangleright g(\lambda_0)$  is optimal
- 12:  $\lambda^* = \lambda_0$
- 13: **break**
- 14: **else if**  $g'(\lambda_0) > 0$  **then**  $\triangleright \lambda^* \geq \lambda_0$
- 15:  $\lambda_-^{(k)} \leftarrow \lambda_0, \lambda_+^{(k)} \leftarrow \lambda_+^{(k-1)}$
- 16: **else**  $\triangleright \lambda^* \leq \lambda_0$
- 17:  $\lambda_+^{(k)} \leftarrow \lambda_0, \lambda_-^{(k)} \leftarrow \lambda_-^{(k-1)}$
- 18: **end if**
- 19:  $k \leftarrow k + 1$   $\triangleright$  Old  $k$  now  $k - 1$
- 20: **until**  $|g(\lambda_+^{(k-1)}) - g(\lambda_-^{(k-1)})| \leq \delta$   $\triangleright$  Up to precision  $\delta$
- 21:  $\lambda^* \leftarrow \lambda_0$
- 22: **end if**
- 23: **return** Optimal mask  $m^* \in \{0, 1\}^D$  with  $\text{supp } m^* = I_{\lambda^*}$

As pointwise infimum over affine functions,  $g$  is concave [22]. A so called *superderivative* of  $g$  can be computed for



every  $\lambda \geq 0$  by differentiating (7) with  $I_\lambda$  fixed, leading to

$$g'(\lambda) = \sum_{i \in I_\lambda} s_i^{(1)} - \kappa. \quad (14)$$

It holds by concavity of  $g$  for  $\lambda \geq 0$

1)  $g'(\lambda) > 0$  (for one superderivative) implies  $\lambda^* \geq \lambda$ ,

2)  $g'(\lambda) < 0$  (for one superderivative) implies  $\lambda^* \leq \lambda$ ,

for all solutions  $\lambda^*$  of  $(\mathcal{D}_1)$ . Using these implications, Algorithm 1 leads to a simple method for approximating  $\lambda^*$  via bisection of intervals. By Theorem 1,  $\lambda^*$  induces a corresponding COPS mask  $m^*$ . For applying Algorithm 1, a proper value  $\lambda_+$  with  $g'(\lambda_+) \leq 0$  has to be found. One possibility is to start with  $\lambda_+ = 1$ . If  $g'(\lambda_+) \geq 0$ ,  $\lambda_+$  can be doubled iteratively. A graphical illustration for Algorithm 1 is given in Fig. 3.

#### E. Convergence and complexity of Algorithm 1

Algorithm 1 has linear order of convergence due to the used bisection method. By (8),  $g(\lambda)$  and  $g'(\lambda)$  can be computed in  $\mathcal{O}(D \cdot \log D)$  by sorting  $\{s_i^{(0)} + \lambda s_i^{(1)}\}$ . This together leads to a deterministic complexity  $\mathcal{O}(D \cdot \log D \cdot \log \frac{1}{\delta})$  of Algorithm 1 for an arbitrary sparsity  $1 \leq \sigma \leq D$  and precision  $\delta > 0$ . Precision  $\delta$  guarantees  $|g(\lambda_{COPS}^*) - g(\lambda^*)| \leq \delta$  for  $\lambda_{COPS}^*$ , the solution found by COPS, and an exact solution  $\lambda^*$  of the dual problem. Using *faster dynamic matrix inverses* (FDMI) [13], the best expected runtime for solving a LP is given by  $\mathcal{O}(D^{2.055} \cdot \log \frac{D}{\delta})$  which is best to our knowledge the fastest method for solving a general LP. With  $D \cdot \log D \ll D^{2.055}$  we see that COPS has drastically smaller complexity than the best standard LP solver FDMI. This speed up is achieved since Algorithm 1 does not invert matrices during optimization, but only sorts scores. The acceleration induced by COPS compared to using standard LP solvers is small in relation to the overall training time for the setup of one-shot pruning before training—usually pruning takes minutes whereas training lasts hours. But if iterative pruning or dynamic sparse training is used, the LP will be solved more often. Thus, COPS’ speed up will be multiplied by the number of score evaluations.

## IV. EXPERIMENTS AND DISCUSSIONS

In this section we present and discuss experimental results for pruning one-shot before training via COPS. To show the superior performance of combining SNIP and GraSP via G-COPS, compared to their individual performance, we evaluate them on the image classification tasks CIFAR-10, CIFAR-100 [12] and Tiny ImageNet [26]. The tested DNN architectures are VGG16 [11], ResNet18 [25] and Wide-ResNet18-2 (WRN18-2) [27]. SNIP and GraSP have shown to be competitive with iterative pruning before training and pruning methods applied during training or after pre-training [7]–[9]. To underline that COPS is not restricted to SNIP/GraSP, we further control SNIP with magnitude scores (M-COPS) and combine SNIP with an unsupervised version of SNIP (USNIP) as a control score, called U-COPS. The U-COPS experiment is conducted on a LeNet-5-Caffe [28], trained on MNIST [28].

Inspired by the interpolation of G-COPS between SNIP and GraSP, shown in Fig. 2, a grid-like search for the multiplier  $\alpha_\kappa$  for  $\kappa_{\min}$  is used. The considered values are  $\alpha_\kappa \in$

$\{0.05, 0.1, 0.2, \dots, 0.9\}$  together with a more fine-grained search close to 1 via  $\alpha_\kappa \in \{0.99, 0.999, 0.9999, 0.99999\}$ . Values  $\alpha_\kappa$  close to one become more interesting for higher pruning rates, as SNIP tends to generate architectures with a weak gradient signal. Thus, a stricter control regarding gradient flow will lead to better results in these cases. We use the same grid-like search also for M-COPS and U-COPS. To achieve fair comparisons, reported  $\alpha_\kappa$  in Tables III and IV are chosen as those with the best mean validation accuracy at early stopping time among all tested  $\alpha_\kappa$ .

#### A. Experimental setup

For the experiments presented in this work we used the deep learning framework PyTorch1.5 [29]. All experiments were run on a single Nvidia GeForce 1080ti GPU.

As common in the literature of pruning before training starts [7]–[9], [20], [21], only weights were pruned. All biases and Batch Normalization parameters were kept trainable. Therefore, the pruning rate  $p$  only considers weights. Like other pruning methods applied before training [7]–[9], [20], [21], also COPS is evaluated on a fixed training setup. Therefore, no hyperparameter optimization was executed, except for  $\alpha_\kappa$ . The training setup and hyperparameters together with references to the used network architectures are given in Table II.

We performed five runs for each reported result for MNIST and CIFAR and three for Tiny ImageNet. Each run was based on a different random seed for weight initialization and ordering of the data. The common data augmentations for CIFAR and Tiny ImageNet were used, with code based on<sup>1</sup>. The training data for MNIST and CIFAR was split randomly 9/1 for training and validation for each random seed. Results are reported for the early stopping epoch, the epoch with the highest validation accuracy.

The considered pruning scores are i.i.d.  $\mathcal{U}(-1, 0)$  random scores, magnitude scores  $-|\Theta_i|$ , GraSP<sup>1</sup> [9], SNIP<sup>2</sup> [7] and USNIP<sup>2</sup> [20]. The determination of SNIP and GraSP scores is discussed in Section III-A. The USNIP method will be introduced in Section IV-D. Except random and magnitude based pruning, all considered pruning methods are data dependent. To obtain a sufficient statistic of the training data, we used the data of 100 randomly chosen training batches to calculate the SNIP, GraSP and USNIP scores. For a pruning method with a single score function, a pruning mask with sparsity  $\sigma$  is computed as a solution to (1). For computing the COPS mask, we implemented Algorithm 1 with precision  $\delta = 10^{-6}$ . To obtain scores in the same range, the target and control scores were normalized to have  $\sum_{i=1}^D |s_i^{(k)}| = 1$ ,  $k = 0, 1$ .

#### B. Experimental results

For CIFAR-10, we tested two different versions of COPS, G-COPS and M-COPS. The results can be seen in Table III. Using GraSP as control score leads to similar results for moderate pruning rates, but better results for higher rates, than controlling the weights’ magnitudes. Therefore, we only

<sup>1</sup><https://github.com/alecwangcq/GraSP>

<sup>2</sup>PyTorch adaptation of <https://github.com/namhoonlee/snip-public>

TABLE II  
EXPERIMENTAL SETUPS WITH REFERENCES TO USED TRAINING AND HYPERPARAMETER SETUP GIVEN IN BRACKETS IN THE TOP ROW.

|                 | MNIST ([7])          | CIFAR-10 ([7])       | CIFAR-100 ([8])   | Tiny ImageNet/ResNet18 ([8]) | Tiny ImageNet/WRN18-2 ([8]) |
|-----------------|----------------------|----------------------|-------------------|------------------------------|-----------------------------|
| #Epochs         | 250                  | 250                  | 160               | 100                          | 100                         |
| Batch Size      | 100                  | 128                  | 128               | 128                          | 128                         |
| Optimizer: SGD- | Momentum 0.9         | Momentum 0.9         | Momentum 0.9      | Momentum 0.9                 | Momentum 0.9                |
| Learning Rate   | 0.1                  | 0.1                  | 0.01              | 0.1                          | 0.1                         |
| LR Decay        | $\times 0.1$         | $\times 0.1$         | $\times 0.2$      | $\times 0.1$                 | $\times 0.1$                |
|                 | every 25k iterations | every 30k iterations | epochs 60/120     | epochs 30/60/80              | epochs 30/60/80             |
| Weight Decay    | $5 \cdot 10^{-4}$    | $5 \cdot 10^{-4}$    | $5 \cdot 10^{-4}$ | $10^{-4}$                    | $10^{-4}$                   |
| Initialization  | Glorot [24]          | He fan_in [25]       | He fan_in [25]    | He fan_in [25]               | He fan_in [25]              |

TABLE III  
EXPERIMENTS ON CIFAR-10 WITH VGG16.

| Method              | $p = 0.9$          | $p = 0.95$    | $p = 0.99$    | $p = 0.995$   |
|---------------------|--------------------|---------------|---------------|---------------|
| Baseline            | $93.41 \pm 0.07\%$ |               |               |               |
| Random              | 91.21%             | 88.83%        | 81.70%        | 66.59%        |
| SNIP                | <u>93.13%</u>      | <u>92.63%</u> | 86.34%        | 10.00%        |
| GraSP               | 92.64%             | 91.94%        | 88.84%        | 76.82%        |
| Magnitude           | 92.94%             | 92.16%        | 10.00%        | 10.00%        |
| G-COPS              | <b>93.15%</b>      | 92.66%        | <b>89.12%</b> | <b>82.67%</b> |
| ( $\alpha_\kappa$ ) | (0.1)              | (0.05)        | (0.99999)     | (0.9999)      |
| M-COPS              | <b>93.15%</b>      | <b>92.68%</b> | 87.71%        | 22.30%        |
| ( $\alpha_\kappa$ ) | (0.1)              | (0.3)         | (0.5)         | (0.4)         |

analyze G-COPS for the bigger datasets CIFAR-100 and Tiny ImageNet, and solely discuss G-COPS in the following paragraphs. Still, controlling magnitudes improves SNIP’s test accuracy by more than 1.3% for  $p = 0.99$ . Remarkably, pruning solely based on magnitudes does not even supply trainable sparse networks for this pruning rate at all.

Tables III and IV compare results for G-COPS, GraSP and SNIP for different pruning rates  $p$ , classification tasks and network architectures. For every task and pruning rate, the best performing G-COPS has better accuracy than SNIP. Using more moderate pruning rates  $p \leq 0.95$ , SNIP and the best G-COPS results lie close together. The gradient flow for SNIP pruned networks is sufficient for these rates, as shown in Fig. 1 for CIFAR-10. Thus, controlling SNIP’s gradient flow only helps marginally. Applying SNIP without constraint seems to be the most economic choice for moderate pruning rates, since  $\alpha_\kappa$  has not to be optimized as hyperparameter. For all models, GraSP and random pruning is outperformed by both, G-COPS and SNIP, for  $p \leq 0.95$ .

As predicted, bigger  $\alpha_\kappa$  achieve the best results for G-COPS with higher pruning rates. For these rates we see G-COPS’ advantage of not being bounded to one pruning score. Using only gradient flow as a score (GraSP) provides better results than using only SNIP scores. Here, not taking the gradient flow into account might result in architectures which do not train at all, like SNIP for  $p = 0.995$  in Table III. But, looking only at the gradient signal can be improved further by additionally optimizing for the sparse architecture with the best influence on the training loss. Even small changes in GraSP’s architecture by a few weights with high saliencies lead to considerably better results, as shown by G-COPS for the ResNet18 on CIFAR-100 and Tiny ImageNet ( $p = 0.99$ ) and the VGG16 on CIFAR-10 ( $p = 0.995$ ). For the latter,

Limited # Classes for Pruning with  $p = 0.995$   
LeNet-5-Caffe on MNIST




Fig. 4. Comparison of SNIP, USNIP and U-COPS. Colored bars represent one standard deviation. Due to unstable training, USNIP has  $89.8 \pm 7.4\%$  accuracy if all classes are used to compute the pruning score which is not shown in the plot for reasons of lucidity.

using G-COPS with  $\alpha_\kappa = 0.9999$ , inducing a similarity of 0.992 between G-COPS and GraSP, surpasses GraSP’s accuracy by 5.85%. Therefore, gradient flow itself should not be the only criterion for choosing a sparse architecture to train. Nevertheless, sufficient gradient flow must be guaranteed for a successful training. As shown, G-COPS offers a possibility to combine these criteria fruitfully.

### C. COPS is not a binary decision between two scores

Tables III and IV show that the grid-like search for  $\alpha_\kappa$  tends to find a G-COPS mask close to either the SNIP or the GraSP mask. The results also show that in almost every case there is a significant performance gap between SNIP and GraSP. Thus, it seems likely that a better performing pruning mask lying in between those two will be close to the superior mask. Such a mask is always found by G-COPS and it clearly improves results compared to GraSP and SNIP, as discussed in Section IV-B. So, even if G-COPS is close to either SNIP or GraSP, it is not a binary decision between those pruning methods. Preferably, COPS is esteemed as a method that combines the better parts of two pruning methods.

TABLE IV  
COMPARISON OF G-COPS ( $m_{\alpha_\kappa}$ ), SNIP ( $M_0$ ) AND GRASP ( $M_1$ );  $\Delta$  DENOTES THE ACCURACY DIFFERENCE TO G-COPS' RESULT.

| Method                          | CIFAR-100/ResNet18 |               |               | Tiny ImageNet/ResNet18 |               |               | Tiny ImageNet/WRN18-2 |               |               |
|---------------------------------|--------------------|---------------|---------------|------------------------|---------------|---------------|-----------------------|---------------|---------------|
|                                 | $p = 0.85$         | $p = 0.95$    | $p = 0.99$    | $p = 0.85$             | $p = 0.95$    | $p = 0.99$    | $p = 0.85$            | $p = 0.95$    | $p = 0.99$    |
| Baseline                        | 74.53 $\pm$ 0.33%  |               |               | 55.97 $\pm$ 0.37%      |               |               | 59.38 $\pm$ 0.25%     |               |               |
| Random                          | 70.71%             | 64.54%        | 50.70%        | 52.59%                 | 46.96%        | 33.87%        | 56.52%                | 52.94%        | 42.63%        |
| SNIP                            | <u>71.61%</u>      | 68.39%        | 56.69%        | <u>56.10%</u>          | <u>54.50%</u> | 37.00%        | <u>58.73%</u>         | <u>57.71%</u> | 51.07%        |
| $\Delta$                        | -0.06%             | -0.13%        | -2.56%        | -0.06%                 | -0.25%        | -7.27%        | -0.03%                | -0.06%        | -1.94%        |
| sim( $M_0, m_{\alpha_\kappa}$ ) | 0.988              | 0.938         | 0.302         | 0.985                  | 0.990         | 0.277         | 0.982                 | 0.981         | 0.271         |
| GraSP                           | 71.16%             | <u>68.40%</u> | <u>58.67%</u> | 54.77%                 | 52.83%        | <u>43.81%</u> | 56.87%                | 57.22%        | <u>52.69%</u> |
| $\Delta$                        | -0.51%             | -0.12%        | -0.58%        | -1.39%                 | -1.92%        | -0.46%        | -1.89%                | -0.55%        | -0.32%        |
| sim( $M_1, m_{\alpha_\kappa}$ ) | 0.351              | 0.350         | 0.994         | 0.343                  | 0.306         | 0.998         | 0.355                 | 0.303         | 0.994         |
| G-COPS                          | <b>71.67%</b>      | <b>68.52%</b> | <b>59.25%</b> | <b>56.16%</b>          | <b>54.75%</b> | <b>44.27%</b> | <b>58.76%</b>         | <b>57.77%</b> | <b>53.01%</b> |
| ( $\alpha_\kappa$ )             | (0.05)             | (0.2)         | (0.9999)      | (0.05)                 | (0.05)        | (0.99999)     | (0.05)                | (0.05)        | (0.9999)      |

#### D. COPS pruning with partial class information

The primal motivation for COPS is given by linking the pruning methods SNIP and GraSP. However, it is not restricted to this combination. In the following, we prune networks with partial class information available at the time of pruning. Meaning,  $T \subset \{1, \dots, C\}$  classes of the original  $C$  classes are used to compute the pruning score. Thus, the classes  $\{1, \dots, C\} \setminus T$  are not known to the pruning procedure. All classes  $\{1, \dots, C\}$  are further used to train and evaluate the pruned network. Analyzing a pruning score in such a way is important if the data used for computing the score will be an unrepresentative sample of the underlying data distribution. As baseline pruning procedure we again use SNIP [7]. The control USNIP score [20] is almost identical to SNIP. They only differ as the loss for USNIP is calculated with labels drawn uniformly at random instead of the real labels. By backpropagation of the loss, every class is included into the calculation of the USNIP score even if the corresponding input image will belong most probably to another class. We evaluate these methods on a small CNN, LeNet-5-Caffe, trained on MNIST. As MNIST is solvable almost perfectly while using only few parameters, we chose the pruning rate  $p = 0.995$ . Results for SNIP, USNIP and U-COPS are shown in Fig. 4.

SNIP pruned networks reach stable results for  $\#T \geq 5$ . However with fewer classes used to calculate the score, SNIP's performance degrades steeply. For these low numbers of regarded classes, USNIP provides better results than SNIP. But, combining both methods via U-COPS reaches the best results for all treated  $\#T$ . In this experiment, U-COPS with a single multiplier  $\alpha_\kappa = 0.9$  reaches better performances than both individual pruning methods for all numbers of considered classes. Also,  $\alpha_\kappa = 0.9$  provides, contrarily to SNIP and USNIP, stable results over all  $\#T$ .

This experiment shows that COPS is not limited to standard pruning methods like SNIP, GraSP or magnitude pruning, but also combines supervised and unsupervised pruning successfully. U-COPS even achieves better results than SNIP for  $\#T = 10$ . Therefore, pruning based on noisy labels via U-COPS, with noise level controlled by  $\alpha_\kappa$ , can generate better trainable sparse architectures. This effect for pruning is comparable to *label smoothing* [30] for standard DNN training.

G-COPS Test Stability for CIFAR-100/ResNet18




Fig. 5. Stability analysis for G-COPS on CIFAR-100 with corresponding test accuracies  $A(\alpha_\kappa)$  displayed on the diagonal. Best accuracy highlighted red.

#### E. Stability of COPS

Fig. 5 analyzes G-COPS' stability for the CIFAR-100 experiment, Table IV, for two pruning rates  $p = 0.85$  and  $p = 0.99$ . On the heatmaps' diagonals, test accuracies for all considered constraint multipliers  $\alpha_\kappa$  are displayed. A cell's color represents the accuracy difference of the G-COPS pruned networks with  $\alpha_\kappa$  in corresponding row and column. As suspected, low constraints reach better results for  $p = 0.85$  than higher ones, whereas for  $p = 0.99$  this behaviour is reversed. The heatmap for  $p = 0.85$  shows that G-COPS is *stable* for all  $\alpha_\kappa$ . Meaning, for small changes of the constraint multiplier  $\alpha_\kappa$ , the corresponding architectures also have small differences in test accuracy. But for  $p = 0.99$ , there is an accuracy gap of 0.9% between  $\alpha_\kappa = 0.9$  and  $\alpha_\kappa = 0.99$ . This performance drop for  $\alpha_\kappa \leq 0.9$  emphasizes the necessity of a strict control on the gradient flow for high pruning rates. Disregarding this accuracy gap, G-COPS also reaches stable results for  $p = 0.99$ .

The heatmap for  $p = 0.85$  also reveals that the best validation accuracy (Table IV:  $\alpha_\kappa = 0.05$ ) not always provides the best test accuracy ( $\alpha_\kappa = 0.2$ ). Still, since accuracy patterns for validation and testing are similar, G-COPS' stability guarantees the best  $\alpha_\kappa$  for validation to be in a neighbourhood of similar test performance as the best  $\alpha_\kappa$  for testing.

#### F. Why solving the dual problem?

Using the strong duality [22] of the problem solved by Algorithm 1 with constraint multiplier  $\alpha_\kappa$ , we see that the

TABLE V

RANDOM SEARCH ON CIFAR-10 WITH VGG16 FOR G-COPS ( $\alpha_\kappa$ ) AND  $\lambda$  FOR THE LINEAR COMBINATION  $S_0 + \lambda S_1$ , 14 HYPERPARAMETERS EACH.

| Method              | Best Test Acc. | Worst Test Acc. | Mean Test Acc. | Std Test Acc. |
|---------------------|----------------|-----------------|----------------|---------------|
| G-COPS              | 93.08%         | 92.79%          | 92.93%         | 0.16%         |
| $S_0 + \lambda S_1$ | 92.77%         | 92.46%          | 92.62%         | 0.20%         |

COPS mask  $m^*$  can equivalently be found as the solution of

$$\min_{m \in \mathcal{X}_{1,\sigma}} S_0(m) + \lambda^* S_1(m), \quad (15)$$

where  $\lambda^*$  solves  $(\mathcal{D}_1)$ . A natural question is why using COPS should be more convenient than combining two scores  $S_0$  and  $S_1$  naively via  $S_0 + \lambda S_1$ . First,  $\alpha_\kappa$  is an interpretable hyperparameter, as shown in Section III-C. An  $\alpha_\kappa$  is closely related to the similarity of the COPS mask and its underlying masks, see Fig. 2. It also controls the strictness of the corresponding constraint  $\alpha_\kappa \cdot \kappa_{\min}$ . For a specific task, this leads to a valuable prior knowledge for finding the right hyperparameter  $\alpha_\kappa$ . *I.e.* high pruning rates result in low gradient flow for SNIP, therefore a strict control of the GraSP score is necessary. Choosing  $\lambda$  as hyperparameter on the other hand only balances the scores which gives no additional insight. Second,  $(0, 1) \ni \alpha_\kappa$  is a much smaller interval than  $(0, \infty) \ni \lambda$ , which improves hyperparameter search. Table V shows a comparison for random search [31] between G-COPS and a pruning mask found by linearly combining  $S_0 + \lambda S_1$ . In order to help out the random search field for  $\lambda \in (0, \infty)$  we used the knowledge provided by G-COPS and further restricted  $\lambda \in (0, 90)$ , as  $80 < \lambda^* < 90$  for the optimal  $\lambda^*$  found by G-COPS for  $\alpha_\kappa = 0.99999$ . We tested 14 different hyperparameters with 5 random seeds each, the same number as for the manually designed grid-like search. The random parameters were chosen uniformly i.i.d. with  $\alpha_\kappa \in (0, 1)$  and  $\lambda \in (0, 90)$ . The results in Table V show that the smaller interval  $(0, 1)$  for  $\alpha_\kappa$ , together with G-COPS’ robustness shown in Section IV-E, clearly speeds up the finding of a well performing  $\alpha_\kappa$ , compared to  $\lambda$  for the linear combination. Naively combining scores linearly does not find an adequate  $\lambda$  in 14 tries. Actually, the best performing hyperparameter found by the linear combination is worse than the worst one tested for G-COPS, showing that G-COPS uses hyperparameters  $\alpha_\kappa$  in a better range than  $\lambda \geq 0$ .

## V. CONCLUSIONS

Training sparse architectures from scratch requires to trade off between different criteria for evaluating the pruned network before training. Both SOTA one-shot pruning methods applied before training starts, SNIP and GraSP, are focused on fulfilling only one of these criteria. On various tasks, network architectures, pruning rates and combinations of pruning scores we have shown that combining two pruning methods via COPS leads to better results than using a single one. Especially for high pruning rates, modifying GraSP only slightly with SNIP can surpass GraSP’s result considerably. However, constraining the gradient flow only improves SNIP marginally for smaller pruning rates. The introduced  $\alpha_\kappa$  is an interpretable

hyperparameter, controlling the strictness of COPS’ constraint. Therefore, COPS is more efficient in finding a well performing hyperparameter than naively combining two scores via a linear combination. COPS defines a pruning mask as a solution of a LP. Algorithm 1 provides a method to compute this mask faster than the best reported expected time for solving a LP for an arbitrary, dense constraint matrix [13]. This might be of intensified interest if more than one pruning step is performed as COPS can also be applied to iterative pruning. But also for structured methods, where pruning is quite restrictive, combining pruning scores raises new opportunities. Likewise, combining more than two pruning scores seems a possibility to improve sparse training even further.

## REFERENCES

- [1] S. Han, J. Pool, J. Tran, and W. Dally, “Learning both weights and connections for efficient neural network,” in *NeurIPS*, 2015.
- [2] Y. LeCun, J. S. Denker, and S. A. Solla, “Optimal brain damage,” in *NeurIPS*, 1990.
- [3] M. C. Mozer and P. Smolensky, “Skeletonization: A technique for trimming the fat from a network via relevance assessment,” in *NeurIPS*, 1989.
- [4] D. W. Blalock, J. J. G. Ortiz, J. Frankle, and J. V. Guttag, “What is the state of neural network pruning?” in *MLSys*, 2020.
- [5] G. Bellec, D. Kappel, W. Maass, and R. Legenstein, “Deep rewiring: Training very sparse deep networks,” in *ICLR*, 2018.
- [6] J. Frankle and M. Carbin, “The lottery ticket hypothesis: Finding sparse, trainable neural networks,” in *ICLR*, 2018.
- [7] N. Lee, T. Ajanthan, and P. H. Torr, “SNIP: Single-shot network pruning based on connection sensitivity,” in *ICLR*, 2019.
- [8] H. Tanaka, D. Kunin, D. L. K. Yamins, and S. Ganguli, “Pruning neural networks without any data by iteratively conserving synaptic flow,” *NeurIPS*, 2020.
- [9] C. Wang, G. Zhang, and R. Grosse, “Picking winning tickets before training by preserving gradient flow,” in *ICLR*, 2020.
- [10] P. Wimmer, J. Mehnert, and A. Condurache, “Freezenet: Full performance by reduced storage costs,” in *ACCV*, 2020.
- [11] K. Simonyan and A. Zisserman, “Very deep convolutional networks for large-scale image recognition,” in *ICLR*, 2015.
- [12] A. Krizhevsky, “Learning multiple layers of features from tiny images,” *University of Toronto*, 2012, <http://www.cs.toronto.edu/~kriz/cifar.html>.
- [13] S. Jiang, Z. Song, O. Weinstein, and H. Zhang, “Faster dynamic matrix inverse for faster lps,” *CoRR*, vol. abs/2004.07470, 2020.
- [14] W. Chen, J. Wilson, S. Tyree, K. Weinberger, and Y. Chen, “Compressing neural networks with the hashing trick,” in *ICML*, 2015.
- [15] J. Xue, J. Li, and Y. Gong, “Restructuring of deep neural network acoustic models with singular value decomposition,” in *Interspeech*, 2013.
- [16] T. N. Sainath, B. Kingsbury, V. Sindhvani, E. Arisoy, and B. Ramabhadran, “Low-rank matrix factorization for deep neural network training with high-dimensional output targets,” in *ICASSP*, 2013.
- [17] H. Li, A. Kadav, I. Durdanovic, H. Samet, and H. P. Graf, “Pruning filters for efficient convnets,” in *ICLR*, 2017.
- [18] S. Han, X. Liu, H. Mao, J. Pu, A. Pedram, M. A. Horowitz, and W. J. Dally, “Eie: Efficient inference engine on compressed deep neural network,” *ACM SIGARCH Computer Architecture News*, vol. 44, no. 3, 2016.
- [19] Y. Chauvin, “A back-propagation algorithm with optimal use of hidden units,” in *NeurIPS*, 1989.
- [20] N. Lee, T. Ajanthan, S. Gould, and P. H. S. Torr, “A signal propagation perspective for pruning neural networks at initialization,” in *ICLR*, 2020.
- [21] S. Verdenius, M. Stol, and P. Forré, “Pruning via iterative ranking of sensitivity statistics,” *CoRR*, vol. abs/2006.00896, 2020.
- [22] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge University Press, 2004.
- [23] G. B. Dantzig, *A History of Scientific Computing*. Association for Computing Machinery, 1990, ch. Origins of the Simplex Method.
- [24] X. Glorot and Y. Bengio, “Understanding the difficulty of training deep feedforward neural networks,” in *AISTATS*, 2010.



- [25] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," *CVPR*, 2016.
- [26] "Tiny imagenet visual recognition challenge," 2015, <http://cs231n.stanford.edu/tiny-imagenet-200.zip>.
- [27] S. Zagoruyko and N. Komodakis, "Wide residual networks," in *BMVC*, 2016.
- [28] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proceedings of the IEEE*, vol. 86, no. 11, 1998.
- [29] A. Paszke *et al.*, "Pytorch: An imperative style, high-performance deep learning library," in *NeurIPS*, 2019.
- [30] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna, "Rethinking the inception architecture for computer vision," in *CVPR*, 2016.
- [31] J. Bergstra and Y. Bengio, "Random search for hyper-parameter optimization," *JMLR*, vol. 13, no. 10, 2012.