

AperTO - Archivio Istituzionale Open Access dell'Università di Torino

SeReNe: Sensitivity-Based Regularization of Neurons for Structured Sparsity in Neural Networks

This is the author's manuscript

Original Citation:

Availability:

This version is available <http://hdl.handle.net/2318/1790716> since 2022-12-20T15:11:10Z

Published version:

DOI:10.1109/TNNLS.2021.3084527

Terms of use:

Open Access

Anyone can freely access the full text of works made available as "Open Access". Works made available under a Creative Commons license can be used according to the terms and conditions of said license. Use of all other works requires consent of the right holder (author or publisher) if not exempted from copyright protection by the applicable law.

(Article begins on next page)

SeReNe: Sensitivity based Regularization of Neurons for Structured Sparsity in Neural Networks

Enzo Tartaglione *Member, IEEE*, Andrea Bragagnolo, Francesco Odierna,
Attilio Fiandrotti, *Senior Member, IEEE*, and Marco Grangetto, *Senior Member, IEEE*

Abstract—Deep neural networks include millions of learnable parameters, making their deployment over resource-constrained devices problematic. SeReNe (Sensitivity-based Regularization of Neurons) is a method for learning sparse topologies with a structure, exploiting *neural sensitivity* as a regularizer. We define the sensitivity of a neuron as the variation of the network output with respect to the variation of the activity of the neuron. The lower the sensitivity of a neuron, the less the network output is perturbed if the neuron output changes. By including the neuron sensitivity in the cost function as a regularization term, we are able to prune neurons with low sensitivity. As entire neurons are pruned rather than single parameters, practical network footprint reduction becomes possible. Our experimental results on multiple network architectures and datasets yield competitive compression ratios with respect to state-of-the-art references.

Index Terms—Sparse networks, regularization, deep networks, pruning, compression.

I. INTRODUCTION

DEEP Neural Networks (DNNs) can solve extremely challenging tasks thanks to complex stacks of (convolutional) layers with thousands of neurons [1]–[3]. Let us define here the *complexity* of a neural network as the number of its learnable parameters: architectures such as AlexNet and VGG have a complexity in the order of 60 and 130 million parameters respectively. Similar architectures are challenging to deploy in scenarios where resources such as the memory or storage are limited. For example, the 8-layers AlexNet [1] memory footprint exceeds 240MB of memory, whereas the 19-layers VGG-Net [2] footprint exceeds 500 MB. The need for compact DNNs is witnessed also by the fact that the Moving Pictures Experts Group (MPEG) of ISO has recently broadened its scope beyond multimedia contents issuing an exploratory call for proposal to compress neural networks [4]. Multiple (complementary) approaches are possible to cope with neural networks memory requirements, inference time and energy consumption:

- Re-designing the network topology. Moving from one architecture to another, possibly forcing a precise neuronal connectivity, or weight sharing, can reduce the number of parameters, or the complexity of the network [3], [5].
- Quantization. Representing the parameters (and activation functions) as fixed-point digits reduces the memory footprint and speeds up computations [6].

- Pruning. Deep architectures need to be over-parametrized [7] to be trained effectively, but redundant parameters can be pruned at inference time [8]–[12]. The present work falls in this latter category.

Pruning techniques aim at learning sparse topologies by selectively dropping synapses between neurons (or neurons altogether when all incoming synapses are dropped). For example, [10] and [11] apply a regularization function promoting low magnitude weights followed by zero thresholding or quantization. Such approaches slash the number of non-zero parameters, allowing to represent the parameters of a layer as a sparse tensor [13]. These methods aim however at pruning parameters independently, so the learned topologies lacks a structure despite sparse. Storing and accessing in memory a randomly sparse matrix entails significant challenges, so it is unclear to which extent such methods could be practically exploited.

This work proposes SeReNe, a method for learning sparse network topologies with a structure, i.e. with fewer neurons altogether. In a nutshell, our method drives *all* the parameters of a neuron towards zero, allowing to prune entire neurons from the network.

First, we introduce the notion *sensitivity of a neuron* as the variation of the network output with respect to the neuron activity. The latter is measured as the post-synaptic potential of the neuron, i.e. the input to the neuron’s activation function. The underlying intuition is that neurons with low sensitivity yield little variation in the network output and thus negligible performance loss if their output changes locally. We also provide computationally efficient bounds to approximate the sensitivity.

Second, we design a regularizer that shrinks *all* parameters of low sensitivity neurons, paving the way to their removal. Indeed, when the sensitivity of a neuron approaches zero, the neuron no longer emits signals and is ready to be pruned.

Third, we propose an iterative two-steps procedure to prune parameters belonging to low sensitivity neurons. Through a cross-validation strategy, we ensure controlled (or even no) performance loss with respect to the original architecture.

Our method allows to learn network topologies which are not only sparse, i.e. with few non-zero parameters, but with fewer neurons (fewer filters for convolutional layers). As a side benefit, smaller and denser architectures may also speedup network execution thanks to a better use of cache locality and memory access pattern.

We experimentally show that SeReNe outperforms state-of-the-art references over multiple learning tasks and network

E. Tartaglione, A. Bragagnolo, F. Odierna, A. Fiandrotti and M. Grangetto are with the Computer Science Department, University of Turin, Torino, ITALY, e-mail: first.last@unito.it

A. Fiandrotti is also with LTCI, Télécom Paris, Institut Polytechnique de Paris, FRANCE, e-mail:attilio.fiandrotti@telecom-paris.fr

architectures. We observe the benefit of structured sparsity when storing the neural network topology and parameters using the *Open Neural Network eXchange* format [14], with a reduction of the memory footprint.

The rest of the paper is structured as follows. In Sec. II we review the relevant literature in neural network pruning; Sec. III provides the definition of sensitivity and practical bounds for its computation and a parameter update rule to “drive” the parameters of low-sensitivity neurons towards zero is presented; Sec. IV presents a practical procedure to prune a network with our scheme and in Sec. V all the empirical results are shown. Finally, in Sec. VI, the conclusions are drawn.

II. RELATED WORK

Approaches towards compact neural networks representations can be categorized in three major groups: altering the network structure, quantizing the parameters and pruning weights. In this section, we review works based on a pruning approach that are most relevant to our work.

In their seminal paper [8], LeCun *et al.* proposed to remove unimportant weights from a network, measuring the importance of each single weight as the increment on the train error when the weight is set to zero. Unfortunately, the complexity of such method would become computationally unbearable in the case of deep topologies with millions of parameters. Due to the scale and the resources required to train and deploy modern deep neural networks, sparse architectures and compression techniques have gained much interest in the deep learning community. Several successful approaches to this problem have been proposed [15]–[18]. While a more in depth analysis on the topic has been published by Gale *et al.* [19], in the rest of this section we provide a summary of the main techniques used to prune deep architectures.

Evolutionary algorithms. Multi-objective sparse feature learning has been proposed by Gong *et al.* [20]: with their evolutionary algorithm, they were able to find a good compromise between sparsity and learning error, at the cost, however, of high computational cost. Similar drawbacks can be found in the work by Lin *et al.*, where convolutional layers are pruned using the artificial bee colony optimization algorithm (dubbed as ABCPruner) [21].

Dropout. Dropout aims at preventing a network from over-fitting by randomly dropping some neurons at learning time [22]. Despite dropout tackles a different problem, it has inspired some techniques aiming at sparsifying deep architectures. Kingma *et al.* [9] have shown that dropout can be seen as a special case of Bayesian regularization. Furthermore, they derive a variational method that allows to use dropout rates adaptively to the data. Molchanov *et al.* [23] exploited such variational dropout to sparsify both fully-connected and convolutional layers. In particular, the parameters having high dropout rate are always ignored and they can be removed from the network. Even if this technique obtains good performance, it is quite complex and it is reported to behave inconsistently when applied to deep architectures [19]. Furthermore, this technique relies on the belief that the Bernoulli probability distribution (to be used with the dropout) is a good variational approximation for the posterior. Another dropout-based

approach is *Targeted Dropout* [24]: here, fine-tuning the ANN model is self-reinforcing its sparsity by stochastically dropping connections. They also target structured sparsity without, however, reaching state-of-the-art performance.

Knowledge distillation. Recently, knowledge distillation [25], [26] received significant attention. The goal in this case is to train a student network (typically shallower or in general having lower complexity) to have the same behavior (in terms of outputs under certain inputs) as a more complex teacher [26] (or as an ensemble of models), reducing the overall computational complexity. Distillation finds application in reducing the prediction of multiple networks into a single one, but can not be applied to minimize the number of neurons for a single network. A recent work is *Few Samples Knowledge Distillation* (FSKD) [27], where a small student network is trained from a larger teacher. In general, in distillation-based techniques, the architecture to be trained is a-priori known, and kept static through all the learning process: in this work, we aim at providing an algorithm which automatically shrinks the deep model’s size with minimal overhead introduced.

Few-shot pruning. Another approach relies on defining the importance of each connection and later remove parameters deemed unnecessary. A recent work by Frankle and Carbin [12] proposed the *lottery ticket hypothesis*, which is having a large impact on the research community. They claim that from an ANN, early in the training, it is possible to extract a sparse sub-network, using a one-shot or iterative fashion: such sparse network, when re-trained, can match the accuracy of the original model. This technique has multiple requirements, like having the history of the training process in order to detect the “lottery winning parameters”, and it is not able to self-tune an automatic thresholding mechanism. Lots of efforts are devoted towards making pruning mechanisms more efficient: for example, Wang *et al.* show that some sparsity is achievable pruning weights at the very beginning of the training process [28], or Lee *et al.*, with their “SNIP”, are able to prune weights in a one-shot fashion [29]. However, these approaches achieve limited sparsity: iterative pruning-based strategy, when compared to one-shot or few-shot approaches, are able to achieve a higher sparsity [30].

Regularization-based pruning. Finally, regularization-based approaches rely on a regularization term (designed to enhance sparsity) to be minimized besides the loss function at training time. Louizos *et al.* propose an ℓ_0 regularization to prune the network parameters during training [31]. Such a technique penalizes non-zero value of a parameter vector, promoting sparse solutions. As a drawback, it requires solving a complex optimization problem, besides the loss minimization strategy and other regularization terms. Han *et al.* propose a multi-step process in which the least relevant parameters are defined, minimizing a target loss function [11]. In particular, it relies on a thresholding heuristics, where all the less important connections are pruned. In [10], a similar approach was followed, introducing a novel regularization term that measures the “sensitivity” of the output wrt. the variation of the parameters. While this technique achieves top-notch sparsity even in deep convolutional architectures, such sparsity is not structured, i.e. the resulting topology includes large




Fig. 1: Representation of the neuron $x_{n,i}$ with activation function $g_{n,i}$.

numbers of neurons with at least one non-zero parameter. Such unstructured sparsity bloats the practically attainable network footprint and leads to irregular memory accesses, jeopardizing execution speedups. In this work we aim at overcoming the above limitations proposing a regularization method that produces a *structured* sparsification, focusing on removing entire neurons instead of single parameters. We also leverage our recent research showing that post-synaptic potential regularization is able to boost generalization over other regularizers [32].

III. SENSITIVITY-BASED REGULARIZATION FOR NEURONS

In this section, we first formulate the sensitivity of a network with respect to the post-synaptic potential of a neuron. Then, we derive a general parameter update rule which relies on the proposed sensitivity term. As reference scenario, a multi-class classification problem with C labels is considered; however, our strategy can be extended to other learning tasks, e.g. regression, in a straightforward way.

A. Preliminaries and Definitions

Let a feed-forward, a-cyclic, multi-layer artificial neural network be composed of $N - 1$ hidden layers. We identify with $n = 0$ the input layer and $n = N$ the output layer, other n values indicate the hidden layers. For the i -th neuron of the n -th layer $x_{n,i}$, we define:

- $y_{n,i}$ as its output,
- \mathbf{y}_{n-1} as its input vector,
- $\theta_{n,i}$ as its own parameters: $w_{n,i}$ the weights and $b_{n,i}$ the bias,

as illustrated in Fig. 1. Each neuron has its own *activation function* $g_{n,i}(\cdot)$ to be applied after some affine function $f_{n,i}(\cdot)$ which can be for example convolution or dot product. Hence, the output of a neuron is given by

$$y_{n,i} = g_{n,i}[p_{n,i}], \quad (1)$$

where $p_{n,i}$ is the *post-synaptic potential* of $x_{n,i}$ defined as:

$$p_{n,i} = f_{n,i}(\theta_{n,i}, \mathbf{y}_{n-1}). \quad (2)$$

B. Neuron Sensitivity

Here we introduce the definition of neuron sensitivity. We recall that we aim at pruning entire neurons rather than single parameters to achieve structured sparsity. Let us we assume that our method is applied to a pre-trained network. To estimate the relevance of neuron $x_{n,i}$ for the task upon which the network was trained, we evaluate the neuron contribution to the network output \mathbf{y}_N . To this end, we first provide an intuition on how small variations of the post-synaptic potential $p_{n,i}$ of neuron $x_{n,i}$ affect the k -th output of the network $y_{N,k}$. By a Taylor series expansion, for small variations of $p_{n,i}$, let us express the variation of $y_{N,k}$ as

$$\Delta y_{N,k} \approx \Delta p_{n,i} \frac{\partial y_{N,k}}{\partial p_{n,i}} \quad (3)$$

where $y_{N,k}$ indicates the k -th output for the output layer. In the case $\Delta y_{N,k} \rightarrow 0, \forall k$, for small variations of $p_{n,i}$, $y_{N,k}$ does not change. Such condition allows to drive the post-synaptic potential $p_{n,i}$ to zero without affecting the network output $y_{N,k}$ (and, for instance, its performance). Otherwise, if $\Delta y_{N,k} \neq 0$, any variation of $p_{n,i}$ might alter the network output, possibly impairing its performance. We can now properly quantify the effect of small changes to the network output by defining the *neuron sensitivity*.

Definition 1: The sensitivity of the network output \mathbf{y}_N with respect to the post-synaptic potential $p_{n,i}$ of neuron $x_{n,i}$ is:

$$S_{n,i}(\mathbf{y}_N, p_{n,i}) = \frac{1}{C} \sum_{k=1}^C \left| \frac{\partial y_{N,k}}{\partial p_{n,i}} \right| \quad (4)$$

where $\mathbf{y}_N \in \mathbb{R}^C$ and $S_{n,i} \in [0; +\infty)$. Intuitively, the higher $S_{n,i}$, the higher the fluctuation of \mathbf{y}_N for small variations of $p_{n,i}$.

Before moving on, we would like to clarify our choice of leveraging the post-synaptic potential $p_{n,i}$ rather than the neuron output $y_{n,i}$ in the equation above. In order to understand our choice, we re-write (4) using the chain rule:

$$S_{n,i}(\mathbf{y}_N, p_{n,i}) = \frac{1}{C} \sum_{k=1}^C \left| \frac{\partial y_{N,k}}{\partial y_{n,i}} \cdot \frac{\partial y_{n,i}}{\partial p_{n,i}} \right|. \quad (5)$$

Without loss of generality, let us assume $\frac{\partial y_{N,k}}{\partial y_{n,i}} \neq 0$ and $g_{n,i}$ corresponds to the well known ReLU activation function. Under the hypothesis that $p_{n,i} < 0$, $\frac{\partial y_{n,i}}{\partial p_{n,i}} = 0$ for the considered ReLU activation. Had we written (4) as a function of the neuron output $y_{n,i}$, the vanishing gradient $\frac{\partial y_{n,i}}{\partial p_{n,i}} = 0$ would have prevented us from estimating the neuron sensitivity. The above consideration applies beyond ReLU to any activation function except for the identity function, for which $y_{n,i} = p_{n,i}$.

C. Bounds on Neuron Sensitivity

Here we provide two computationally-efficient bounds to the sensitivity function above that can be practically exploited. Popular frameworks for DNN training rely on differentiation frameworks such as *autograd*, for automatic variable differentiation along computational graphs. Such frameworks take as input some objective function J and automatically compute all

the gradients along the computational graph. In order to get $S_{n,i}$ as an outcome from the differentiation engine, we define

$$S_{n,i}(\mathbf{y}_N, p_{n,i}) = \frac{\partial J}{\partial p_{n,i}} \quad (6)$$

where J is a proper function. In Appendix A we show that such function turns to be:

$$J = \frac{1}{C} \sum_{k=1}^C \int \left| \frac{\partial y_{N,k}}{\partial p_{n,i}} \right| dp_{n,i} \quad (7)$$

Therefore, computing the sensitivity in (4) requires C calls to the differentiation engine. In the following with some little algebra we derive a lower and upper bound to Def. 1 that we show to be particularly useful from a computational perspective.

Let the objective function to differentiate be

$$J^l = \frac{1}{C} \sum_{k=1}^C y_{N,k}. \quad (8)$$

The automatic differentiation engine called on S^l will return

$$\frac{\partial J^l}{\partial p_{n,i}} = \frac{1}{C} \sum_{k=1}^C \frac{\partial y_{N,k}}{\partial p_{n,i}} \quad (9)$$

According to the triangular inequality, a lower bound to the sensitivity in (4) can be computed as

$$S_{n,i}^l = \frac{1}{C} \left| \sum_{k=1}^C \frac{\partial y_{N,k}}{\partial p_{n,i}} \right| \leq \frac{1}{C} \sum_{k=1}^C \left| \frac{\partial y_{N,k}}{\partial p_{n,i}} \right| \quad (10)$$

$S_{n,i}^l$ can be conveniently evaluated differentiating over (8) (and taking the absolute value) with a single call to the differentiation engine. As shown in (10), this gives us a lower bound estimation over the neuron sensitivity.

In order to estimate an upper bound to $S_{n,i}$, we rewrite (4) as

$$S_{n,i} = \frac{1}{C} \sum_{k=1}^C \left| \frac{\partial y_{N,k}}{\partial \mathbf{y}_{N-1}} \cdot \prod_{l=n+1}^{N-1} \frac{\partial \mathbf{y}_l}{\partial \mathbf{y}_{l-1}} \cdot \delta_{n,i} \frac{\partial y_{N,k}}{\partial p_{n,i}} \right| \quad (11)$$

However, $\forall k$ we have in common the term

$$\begin{aligned} \Gamma_{n,i} &= \prod_{l=n+1}^{N-1} \frac{\partial \mathbf{y}_l}{\partial \mathbf{y}_{l-1}} \cdot \delta_{n,i} \frac{\partial y_{N,i}}{\partial p_{n,i}} \\ &\leq \prod_{l=n+1}^{N-1} \left| \frac{\partial \mathbf{y}_l}{\partial \mathbf{y}_{l-1}} \right| \cdot \delta_{n,i} \left| \frac{\partial y_{N,i}}{\partial p_{n,i}} \right| = \Gamma_{n,i}^u \end{aligned} \quad (12)$$

where $\delta_{n,i}$ is a one-hot vector selecting the i -th neuron at the n -th layer and $|\cdot|$ is an element-wise operator. Hence, we rewrite (11) as

$$S_{n,i}^u = \frac{1}{C} \left(\sum_{k=1}^C \left| \frac{\partial y_{N,k}}{\partial \mathbf{y}_{N-1}} \right| \right) \cdot \Gamma_{n,i}^u \geq S_{n,i}. \quad (13)$$

Thus, we have shown that $S_{n,i}^u$ is an upper bound to the sensitivity in (4). Upper and lower bounds are here obtained for two main reasons: computational efficiency and relaxing/tightening conditions on the sensitivity itself. We will see in Sec. V-A a typical population distribution of the sensitivities on a pre-trained network, comparing (4), (10) and (13).

In the following, we exploit the formulation of the the Sensitivity function (4) and its two bounds (10), (13) to define a parameter update rule.

D. Parameters Update Rule

Now we show how the proposed sensitivity definition can be exploited to promote neuron sparsification. As hinted before, if the sensitivity $S_{n,i}$ of neuron $x_{n,i}$ is small, i.e $S_{n,i} \rightarrow 0$, then neuron $x_{n,i}$ yields a small contribution to the i -th network output $y_{N,i}$ and its parameters may be moved towards zero with little perturbation to the network output. To this end, we define the *insensitivity* function $\bar{S}_{n,i}$ as

$$\bar{S}_{n,i} = \max\{0, 1 - S_{n,i}\} = (1 - S_{n,i}) \cdot \Theta(1 - S_{n,i}) \quad (14)$$

where $\Theta(\cdot)$ is the one-step function. The higher the insensitivity of neuron $x_{n,i}$ (i.e., $\bar{S}_{n,i} \rightarrow 1$ or equivalently $S_{n,i} \rightarrow 0$), the less the neuron affects the network output. Therefore, if $\bar{S}_{n,i} \rightarrow 1$, then neuron $x_{n,i}$ contributes little to the network output and its parameters $w_{n,i,j}$ can be driven towards zero without significantly perturbing the network output. Using the insensitivity definition in (14), we propose the following update rule:

$$w_{n,i,j}^{t+1} = w_{n,i,j}^t - \eta \frac{\partial L}{\partial w_{n,i,j}^t} - \lambda w_{n,i,j}^t \bar{S}_{n,i} \quad (15)$$

where

- the first contribution term is the classical minimization of a loss function L , ensuring that the network still solves the target task, e.g. classification;
- the second one represents a penalty applied to the parameter $w_{n,i,j}$ belonging to the neuron $x_{n,i}$ which is proportional to the insensitivity of the output to its variations.

Finally, since

$$\frac{\partial p_{n,i}}{\partial y_{n-1,j}} = w_{n,i,j} \quad (16)$$

we rewrite (15) as

$$w_{n,i,j}^{t+1} = w_{n,i,j}^t - \eta \frac{\partial L}{\partial w_{n,i,j}^t} - \lambda \dot{S}_{n,i,j} \quad (17)$$

where

$$\dot{S}_{n,i,j} = \left[w_{n,i,j} - \frac{\text{sign}(w_{n,i,j})}{C} \sum_{k=1}^C \left| \frac{\partial y_{N,k}}{\partial y_{n-1,j}} \right| \right] \cdot \Theta(1 - S_{n,i}) \quad (18)$$

A step-by-step derivation is provided in Appendix D. From (18) we can better understand the effect of the proposed penalty term: as expected by our discussion above, $\dot{S}_{n,i,j}$ is inversely proportional to the impact on the output for variations of the input for the neuron $x_{n,i}$.

E. Local neuron sensitivity-based regularization

We propose now an approximate formulation of the sensitivity function in (4) based only on the post-synaptic potential and output of a neuron that we will refer to as the *local* sensitivity. Let us recall that for each neuron $x_{n,i}$ the sensitivity provided by Definition 1 measures the overall impact of a




Fig. 2: High-level view of the SeReNe procedure.

given neuron $x_{n,i}$ on the network output taking into account all the following neurons involved in the computation.

Definition 2: The *local* neuron sensitivity of the output $y_{n,i}$ with respect to the post-synaptic potential $p_{n,i}$ of the neuron $x_{n,i}$ is defined as:

$$\tilde{S}_{n,i} = \left| \frac{\partial y_{n,i}}{\partial p_{n,i}} \right| \quad (19)$$

In the case of ReLU-activated networks, it simply reads

$$\tilde{S}_{n,i} = \Theta(p_{n,i}) \quad (20)$$

Under this setting, the update rule (17) simplifies to

$$w_{n,i,j}^{t+1} = w_{n,i,j}^t - \eta \frac{\partial L}{\partial w_{n,i,j}^t} - \lambda w_{n,i,j}^t \Theta(-p_{n,i}), \quad (21)$$

i.e., the penalty is applied only in case the neuron stays off. While local sensitivity is a looser approximation of (1), it is far less complex to compute especially for ReLU-activated neurons.

IV. THE SERENE PROCEDURE

This section introduces a practical procedure to prune neurons from a neural network \mathcal{N} leveraging the sensitivity-based regularizer introduced above.

Let us assume \mathcal{N} has been preliminary trained at some task over the dataset D achieving performance (e.g., classification accuracy) A . We do not put any constraint over the actual training method, training set or network architecture. Alg. 1 summarizes the procedure in pseudo-code. In a nutshell, the procedure consists in iteratively looping over the *Regularization* and *Thresholding* procedures.

At the beginning of the loop, dataset D is split into disjoint subset V (used for validation purposes) and U (to update the

network). At line 5, the regularization procedure (summarized in Alg. 2) trains \mathcal{N} over D according to (15) driving towards zero parameters of neurons with low sensitivity. The loop ends if the performance of the regularized network falls below threshold A . Otherwise, the thresholding procedure sets to zero parameters below threshold T and prunes neurons such that all parameters are equal to zero. The output of the procedure is the pruned network, i.e. with fewer neurons, \mathcal{N}^* . The Regularization and Thresholding procedures are detailed in the following. A graphical high-level representation of SeReNe is also displayed in Fig. 2.

Algorithm 1 The SeReNe procedure

Input: Trained network \mathcal{N} , dataset D ,

Target performance A , PWE , TWT

Output: Pruned network \mathcal{N}^*

```

1: procedure SERENE( $\mathcal{N}, D, A, PWE, TWT$ )
2:    $\mathcal{N}^* \leftarrow \mathcal{N}$ 
3:   while true do
4:      $U, V \leftarrow \text{RANDOMSPPLIT}(D)$ 
5:      $\mathcal{N} \leftarrow \text{REGULARIZATION}(\mathcal{N}, U, V, PWE)$ 
6:     if PERFORMANCE( $\mathcal{N}, V$ ) <  $A$  then
7:       break
8:      $\mathcal{N}^* \leftarrow \mathcal{N}$ 
9:      $\mathcal{N} \leftarrow \text{THRESHOLDING}(\mathcal{N}, V, TWT)$ 
return  $\mathcal{N}^*$ 

```

A. Regularization

This procedure takes in input a network \mathcal{N} and returns a regularized network according to the update rule (15). Namely, the procedure iteratively trains \mathcal{N} on U and validates it on V for multiple epochs. Let \mathcal{N}^r represent the *best* regularized network found at a given time according to the loss function. For each iteration, the procedure operates as follows.

First (line 5), \mathcal{N} is trained for one epoch over U : the results is a regularized network according to (15).

Second (line 6), this network is validated on V . If the loss is lower than the loss of \mathcal{N}^r over V , then \mathcal{N} takes the place of \mathcal{N}^r (line 7). If \mathcal{N}^r is not updated for PWE (*Plateau Waiting Epochs*) epochs, we assume we have reached a performance plateau. In this case, the procedure ends and returns the sensitivity-regularized network \mathcal{N}^r .

Algorithm 2 The regularization procedure

Input: Model \mathcal{N} , data sets V and U , PWE

Output: The sensitivity-regularized network \mathcal{N}^r

```

1: procedure REGULARIZATION( $\mathcal{N}, U, V, PWE$ )
2:    $\mathcal{N}^r \leftarrow \mathcal{N}$   $\triangleright \mathcal{N}^r$  is best regularized network on  $V$ 
3:    $epochs \leftarrow 0$ 
4:   while  $epochs < PWE$  do
5:      $\mathcal{N} \leftarrow \text{TRAIN}(\mathcal{N}, U)$   $\triangleright 1$  train epoch on  $U$ 
6:      $epochs ++$ 
7:     if LOSS( $\mathcal{N}, V$ ) < LOSS( $\mathcal{N}^r, V$ ) then
8:        $\mathcal{N}^r \leftarrow \mathcal{N}$ 
9:      $epochs \leftarrow 0$ 
return  $\mathcal{N}^r$ 

```

B. Thresholding

The *thresholding* procedure is where the parameters of neurons with low sensitivity are thresholded to zero. Namely, parameters whose absolute value is below threshold T are pruned as

$$w_{n,i,j} = \begin{cases} w_{n,i,j} & |w_{n,i,j}| > T \\ 0 & \text{otherwise.} \end{cases} \quad (22)$$

The pruning threshold T is selected so that the performance (or, in other words, the loss on V) worsens at most of a relative value we call *thresholding worsening tolerance (TWT)* we provide as hyper-parameter. We expect the loss function to be locally a smooth, monotone function of T , for small values of T . The threshold T can be found using linear search-based heuristics. We can however reduce this using a bisection approach, converging to the optimal T value in log-time steps. Because of the stochasticity introduced by mini-batch based optimizers, parameters pruned during a thresholding iteration may be reintroduced by the following regularization iteration. In order to overcome this effect, we enforce that pruned parameters can no longer be updated during the following regularizations (we term this behavior as *parameter pinning*). To this end, the update rule (15) is modified as follows:

$$w_{n,i,j}^{t+1} = \begin{cases} w_{n,i,j}^t - \eta \frac{\partial L}{\partial w_{n,i,j}^t} - \lambda w_{n,i,j}^t \bar{S}_{n,i} & w_{n,i,j}^t \neq 0 \\ 0 & w_{n,i,j}^t = 0 \end{cases} \quad (23)$$

We have noticed that without parameter pinning, the compression of the network may remain low because the noisy gradient estimates in a mini-batch that keep reintroducing previously pruned parameters. On the contrary, by adding (23) a lower number of epochs are sufficient to achieve much higher compression.

V. RESULTS

In this section we experiment with our proposed neuron pruning method comparing the four sensitivity formulations we introduced in the previous section:

- SeReNe (exact) - the exact formulation in (4);
- SeReNe (LB) - the lower bound in (10);
- SeReNe (UB) - the upper bound in (13);
- SeReNe (local) - the local version in (19);
- ℓ_2 + pruning - is a baseline reference where we replace our sensitivity-based regularization term with a standard ℓ_2 term (all the rest of the framework is identical).

We experiment over different combinations of architectures and datasets commonly used as benchmarks in the relevant literature:

- LeNet-300 on MNIST (Table I and Table II),
- LeNet-5 on MNIST (Table III),
- LeNet-5 on Fashion-MNIST (Table IV),
- VGG-16 on CIFAR-10 (Table V and Table VI),
- ResNet-32 on CIFAR-10 (Table VII),
- AlexNet on CIFAR-100 (Table VIII),
- ResNet-101 on ImageNet (Table IX).

Notice that the VGG-16, AlexNet and ResNet-32 architectures are modified to fit the target classification task (CIFAR-10 and




Fig. 3: Population of sensitivities S and relative lower S^l and upper S^u bounds for a LeNet-5 architecture pre-trained on MNIST. Vertical bars indicate relative mean values.

CIFAR-100). The validation set (V) size for all experiments is 10% of the training set.

The pruning performance is evaluated according to multiple metrics.

- The *compression ratio* as the ratio between the number of parameters in the original network and the number of remaining parameters after pruning (the higher the better).
- The number of remaining neurons (or filters for convolutional layers) after pruning.
- The size of the networks when stored on disk in the popular ONNX format [14] (`.onnx` column). ONNX files are then lossless compressed using the LempelZivMarkov algorithm (LZMA) [35] (`.7z` column).

In our experiments, we compare with all available references for each combination of architecture and dataset. For this reason, the reference set may vary from experiment to experiment. Our algorithms are implemented in Python, using PyTorch 1.5, and simulations are run on a RTX2080 NVIDIA GPU with 8GB of memory.¹

A. Preliminary experiment

To start with, we plot the sensitivity distribution for a LeNet-5 network trained on the MNIST dataset (SGD with learning rate $\eta = 0.1$ weight-decay 10^{-4}). This network will also be used as baseline in Sec. V-C. Fig. 3 shows SeReNe (exact) (red), SeReNe (LB) (green) and SeReNe (UB) (blue); the vertical bars represent the mean values. As expected, SeReNe (LB) and SeReNe (UB) under-estimate and over-estimate SeReNe (exact), respectively. Interestingly, SeReNe (UB) sensitivity values lie in the range $[10^{-4}; 10^{-2}]$ while both for SeReNe (exact) and SeReNe (LB) show a longer trail towards smaller figures, whereas all distributions look similar. In the following, we will experimentally evaluate the three sensitivity formulations in terms of pruning effectiveness.

¹The source code is available at <https://github.com/EIDOSlab/SeReNe.git>

TABLE I: LeNet-300 trained on MNIST (1.65% error rate).

Approach	Remaining parameters (%)			Compr. ratio	Remaining neurons	Network size [kB]		Training time (s/epoch)	Top-1 (%)
	FC1	FC2	FC3			.onnx	.7z		
Baseline	100	100	100	1x	[300]-[100]-[10]	1043	→ 933	3.65	1.44
Han <i>et al.</i> [11]	8	9	26	12.2x	-	-	-	-	1.6
Tartaglione <i>et al.</i> [10]	2.25	11.93	69.3	27.87x	[251]-[88]-[10]	-	-	-	1.65
ℓ_2 +pruning	2.44	15.76	68.50	23.26x	[212]-[82]-[10]	723	→ 64	3.65	1.66
SeReNe (exact)	1.42	9.54	60.9	42.55x	[159]-[75]-[10]	538	→ 46	13.25	1.64
SeReNe (UB)	22.45	60.81	87.75	3.71x	[295]-[92]-[10]	1016	→ 324	5.13	1.67
SeReNe (LB)	1.51	10.05	60.53	39.79x	[164]-[78]-[10]	557	→ 55	4.88	1.65
SeReNe (local)	3.85	32.53	73.49	13.81x	[251]-[86]-[10]	859	→ 119	3.83	1.64

TABLE II: LeNet-300 trained on MNIST (1.95% error rate).

Approach	Remaining parameters (%)			Compr. ratio	Remaining neurons	Network size [kB]		Training time (s/epoch)	Top-1 (%)
	FC1	FC2	FC3			.onnx	.7z		
Baseline	100	100	100	1x	[300]-[100]-[10]	1043	→ 933	3.65	1.44
Sparse VD [23]	-	-	-	68x	-	-	-	-	1.92
SWS [33]	-	-	-	23x	-	-	-	-	1.94
Tartaglione <i>et al.</i> [10]	0.93	1.12	5.9	103x	[179]-[88]-[10]	-	-	-	1.95
DNS [34]	1.8	1.8	5.5	56x	-	-	-	-	1.99
ℓ_2 +pruning	1.22	8.77	61.10	41.95x	[167]-[76]-[10]	566	→ 42	3.65	1.97
SeReNe (exact)	0.76	5.85	49.77	66.28x	[148]-[70]-[10]	498	→ 38	13.25	1.93
SeReNe (UB)	13.67	50.76	84.47	5.47x	[293]-[91]-[10]	1008	→ 240	5.13	1.95
SeReNe (LB)	0.75	5.79	49.3	66.41x	[146]-[70]-[10]	492	→ 37	4.88	1.95
SeReNe (local)	1.7	19.94	63.59	25.07x	[192]-[83]-[10]	656	→ 70	3.83	1.93

B. LeNet300 on MNIST

As a first experiment, we prune a LeNet-300 architecture, which consists of three fully-connected layers with 300, 100 and 10 neurons, respectively trained over the MNIST dataset. We pre-trained LeNet-300 via SGD with learning rate $\eta = 0.1$ and $PWE = 20$ epochs with $\lambda = 10^{-5}$, $TWT = 0.3$ for SeReNe (exact), SeReNe (LB) SeReNe (UB) and $\lambda = 10^{-5}$, $TWT = 1$ for SeReNe (local). The related literature reports

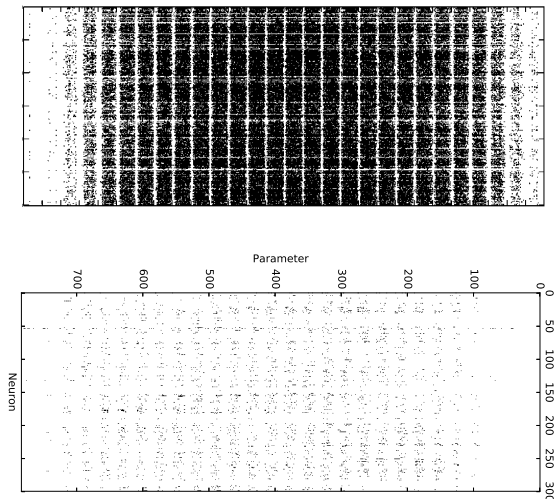


Fig. 4: Parameters distribution in FC1 of LeNet-300 trained on MNIST from Han *et al.* [11] (top) and the proposed SeReNe (bottom). In black the remaining parameters.

mainly i) results for classification errors around 1.65% (Table I) and ii) results for errors in the order of 1.95% (Table II). For this reason, we trained for about 1k epochs to achieve 1.95% error rate and for additional 2k epochs to score a 1.65% error rate.

SeReNe outperforms the other methods leads both in terms of compression ratio and number of pruned neurons. SeReNe (exact) achieves a compression ratio of 42.55 \times and the number of remaining neurons in the hidden layers drops from 300 to 159 and from 100 to 75 respectively. SeReNe (LB) enjoys comparable performance with respect to SeReNe (exact) despite lower computational cost (see below).

For the 1.95% error band (Tab. II), SeReNe (LB) performs is more effective at pruning parameters than SeReNe (exact), allowing lower error. In this case, SeReNe (LB) achieves comparable compression to other state-of-the-art approaches, except for Tartaglione *et al.* [10]. In this case, however, when we compare the final architecture, we see that [10] prunes less neurons than SeReNe, despite a higher compression ratio. Evidently, [10] enhances unstructured sparsity, while SeReNe exploits structured sparsity, resulting in more entirely-removed neurons. SeReNe (LB) prunes more parameters than SeReNe (UB), we hypothesize because (13) overestimates the sensitivity of the parameters and prevents them to be pruned. On the other side, SeReNe (LB) underestimates the sensitivity, however small λ values sets this off. SeReNe (local) prunes less parameters than the other SeReNe formulations as it relies on a locally computed sensitivity formulation despite lower complexity. Concerning training time (second column from the right), SeReNe (local) is fastest and introduces very

TABLE III: LeNet-5 trained on MNIST.

Approach	Remaining parameters (%)				Compr. ratio	Neurons	Network size [kB]		Top-1 (%)
	Conv1	Conv2	FC1	FC2			.onnx	.7z	
Baseline	100	100	100	100	1x	[20]-[50]-[500]-[10]	1686	→ 1510	0.68
Sparse VD [23]	33	2	0.2	5	280x	-	-	-	0.75
Han <i>et al.</i> [11]	66	12	8	19	11.9x	-	-	-	0.77
SWS [33]	-	-	-	-	162x	-	-	-	0.97
Tartaglione <i>et al.</i> [10]	67.6	11.8	0.9	31.0	51.1x	[20]-[48]-[344]-[10]	-	-	0.78
DNS [34]	14	3	0.7	4	111x	-	-	-	0.91
ℓ_2 +pruning	60.20	7.37	0.61	22.14	72.3	[19]-[37]-[214]-[10]	577	→ 46	0.8
SeReNe (LB)	33.75	3.25	0.27	10.22	177.05x	[11]-[26]-[113]-[10]	208	→ 19	0.8

TABLE IV: LeNet-5 trained on Fashion-MNIST.

Approach	Remaining parameters (%)				Compr. ratio	Neurons	Network size [kB]		Top-1 (%)
	Conv1	Conv2	FC1	FC2			.onnx	.7z	
Baseline	100	100	100	100	1x	[20]-[50]-[500]-[10]	1686	→ 1510	8.1
Sparse VD [23]	-	-	-	-	6.98x	-	-	-	8.53
Tartaglione <i>et al.</i> [10]	76.2	32.56	6.5	44.02	11.74x	[20]-[47]-[470]-[10]	-	-	8.5
ℓ_2 +pruning	85.80	34.13	4.57	55.24	14.36x	[20]-[50]-[500]-[10]	1496	→ 197	8.44
SeReNe (LB)	85.71	32.14	3.63	52.03	17.04x	[20]-[49]-[449]-[10]	1494	→ 46	8.47

little computational overhead, SeReNe (UB) and SeReNe (LB) have comparable training times and the slowest is the SeReNe (exact), approximately 2.7x slower than its boundaries. In the light of the good trade-off between ability to prune neurons, error rate and training time of SeReNe (LB), in the following we will restrict our experiments to this sensitivity formulation.

Fig. 4 (bottom) shows the location of the parameters not pruned by SeReNe (exact) in LeNet300 first fully-connected layer (black dots). For comparison, we report the equivalent image from Fig. 4 of [11] (top). Our method yields completely blank columns in the matrix that can be represented in memory as uninterrupted sequences of zeroes. When stored on disk, LZMA compression (.7z column) is particularly effective at encoding long sequences of the same symbol, which explains the 10x compression rate it achieves (from 538 to 46 kB) over the .onnx file.

Finally, we perform an ablation study to assess the impact of a simpler ℓ_2 -only regularization, i.e. classical weight decay, in place of our sensitivity-based regularizer. Towards this end, we retrain LeNet-300 with $\lambda = 0$ and a weight-decay set to 10^{-4} in its place (line ℓ_2 +pruning in the tables above). We point out in (15) that the sensitivity can be interpreted as a weighting factor for the ℓ_2 -regularization. Using weight-decay is equivalent to assuming all the parameters have the same sensitivity. For this experiment, we used $\eta = 0.1$, $PWE = 5$ and $TWT = 0$ ($TWT > 0$ significantly and uncontrollably worsens the performance). Table I shows that such method is less effective at pruning neurons than SeReNe (LB), which removes 15% more neurons. Similar conclusions can be drawn also if higher error is tolerated, as in Table II. The ℓ_2 +pruning has been performed for comparison in all following experiments in the paper yielding the same results.

C. LeNet5 on MNIST

Next, we repeat the previous experiment over the LeNet-5 [36] architecture, preliminarily trained as for the LeNet-300 above, yet with SGD with learning rate $\eta = 0.1$ and $PWE = 20$ epochs. We experiment with SeReNe (LB) with parameters ($\lambda = 10^{-4}$, $TWT = 1.45$). For this architecture, our method requires about 500 epochs to achieve the same error range as other state of the art references. According to Table III, SeReNe (LB) approaches the classification accuracy of its competitors outperforms the considered references in terms of compression ratio and pruned neurons.

In this case, the benefits coming from the structured sparsity are evident: the uncompressed network storage footprint decreases from 1686 kB to 208 kB (-90%), which after lossless compression further decreases to 19 kB with a 0.12% performance drop only.

D. LeNet5 on Fashion-MNIST

Then, we experiment with the same LeNet-5 architecture on the Fashion-MNIST [37] dataset. Fashion-MNIST has the same size of the MNIST dataset, yet it contains natural images of dresses, shoes, etc. and so it is harder to classify than MNIST since the images are not sparse as MNIST digits. In this experiment we used SGD with learning rate $\eta = 0.1$ and $PWE = 20$ epochs. For SeReNe (LB) we used $\lambda = 10^{-5}$ and $TWT = 1$ for about 2k epochs.

Unsurprisingly, the average compression ratio is lower than for MNIST: since the classification problem is much harder than MNIST (Sec. V-C), more complexity is required and SeReNe, in order not to degrade the Top-1 performance, is not pruning as much as it did for the MNIST experiment. Most importantly, the SeReNe (LB) compressed network is 46 kB only, despite the higher number of pruned parameters.

TABLE V: VGG-like architecture with 1 fully connected layer (*VGG-1*) trained on CIFAR-10.

Approach	Remaining parameters (%) [neurons]						FC1	Compr. ratio	Network size [MB]		Top-1 (%)
	Conv1	Conv2	Conv3	Conv4	Conv5	.onnx			.7z		
Baseline	100	100	100	100	100	-	1x	57.57	→	51.51	7.36
	[64]	[128]	[256]	[512]	[512]	[10]					
	[64]	[128]	[256]	[512]	[512]						
ℓ_2 +pruning	11.86	15.07	6.59	0.36	0.11	66.70	88.84x	13.58	→	1.14	7.79
	[23]	[126]	[250]	[406]	[60]	[10]					
	[64]	[123]	[251]	[108]	[81]						
SeReNe (LB)			[250]	[128]	[398]						
	10.18	11.68	4.73	0.20	0.05	61.11	124.82x	11.56	→	0.97	7.8
	[23]	[126]	[250]	[382]	[65]	[10]					
	[64]	[123]	[251]	[93]	[76]						
			[250]	[136]	[373]						

TABLE VI: VGG-like architecture with 2 fully connected layers (*VGG-2*) trained on CIFAR-10.

Approach	Remaining parameters (%) [neurons]							Compr. ratio	Network size [MB]		Top-1 (%)		
	Conv1	Conv2	Conv3	Conv4	Conv5	FC1	FC2		.onnx	.7z			
Baseline	100	100	100	100	100	100	100	1x	58.61	→	52.44	6.16	
	[64]	[128]	[256]	[512]	[512]	[512]	[10]						
	[64]	[128]	[256]	[512]	[512]								
Sparse-VD [23]	-	-	-	-	-	-	-	48x	-	-	-	7.3	
	ℓ_2 +pruning	27.62	30.74	13.67	0.88	0.24	1.88	70.78	40.96x	34.42	→	2.86	7.21
	[44]	[126]	[247]	[498]	[409]	[367]	[10]						
SeReNe (LB)	[60]	[120]	[247]	[463]	[417]								
			[243]	[79]	[461]								
	25.9	26.38	9.75	0.48	0.15	1.24	70	57.99x	29.41	→	2.47	7.25	
	[44]	[126]	[247]	[498]	[354]	[367]	[10]						
	[60]	[120]	[247]	[433]	[366]								
			[243]	[65]	[459]								

TABLE VII: ResNet-32 trained on CIFAR-10.

Approach	Remaining parameters (%) [neurons]					FC1	Compr. ratio	Network size [MB]		Top-1 (%)	
	Conv1	Block1	Block2	Block3	Block3			.onnx	.7z		
Baseline	100	100	100	100	100	100	1x	1.84	→	1.63	7.36
	[64]	[160]	[320]	[640]	[10]						
Sparse VD [23]	-	-	-	-	-	-	2.5x	-	-	-	8.16
	ℓ_2 +pruning	65.97	33.30	33.41	26.32	88.75	3.51x	1.82	→	0.54	8.08
SeReNe (LB)	[14]	[157]	[319]	[633]	[10]						
	60.19	24.52	24.14	17.84	81.88	5.03x	0.87	→	0.37	8.09	
	[12]	[93]	[203]	[364]	[10]						

E. VGG on CIFAR-10.

Next, we experiment with two popular implementations of the VGG architecture [2]. We recall that VGG consists in 13 convolutional layers arranged in 5 groups of, respectively, 2, 2, 3, 3, 3 layers, with 64, 128, 256, 512, 512 filters per layer respectively. *VGG-1* is a VGG implementation popular in CIFAR-10 experiments that includes only one fully-connected layer as output layer and is pre-trained on ImageNet.² *VGG-2* [23] is similar to VGG-1 but includes one hidden fully connected layer with 512 neurons before the output layer. We experiment over the CIFAR-10 dataset, which consists of

50k 32×32 , RGB images for training and 10k for testing, distributed in 10 classes. For both VGG-1 and VGG-2 we have used SGD with learning rate $\eta = 0.01$ and $PWE = 20$ epochs. For the SeReNe (LB), we used $\lambda = 10^{-6}$ and $TWT = 1.5$. Both architectures were pruned for approximately 1k epochs and Tables V and VI detail the pruned topologies. For each architecture, we detail the number of surviving filters (convolutional layers) or neurons (fully connected layers) for each layer within square brackets. The tables show that SeReNe introduces a significantly structured sparsity for both VGG-1 and VGG-2 and outperforms Sparse-VD [23] in terms of compression ratio. We are able to prune a significant number

²<https://github.com/kuangliu/pytorch-cifar>

of filters also in the convolutional layers; as an example, the 3 layers in block Conv4 are reduced to [382]-[93]-[136] for VGG-1 and [498]-[433]-[65] for VGG-2. This has a positive impact on the networks footprint: VGG-1 memory footprint drops from 57.57 MB to 11.56 MB for the pruned network, while the *7zip* compressed representation is 0.97 MB only. For VGG-2, the memory footprint drops from 58.61 MB to 29.41 MB, while the compressed file representation amounts to 2.47 MB.

F. ResNet-32 on CIFAR-10

We then evaluate SeReNe over the ResNet-32 architecture [3] trained on the CIFAR-10 dataset using SGD with learning rate $\eta = 0.001$, momentum 0.9, $\lambda = 10^{-5}$, $TWT = 0$ and $PWE = 10$. Table VII shows the resulting architecture. Due to the number of layers, we represent the network architecture in five different blocks: the first correspond to the first convolutional layer that takes in input the original input image, the last represent the fully-connected output layer. The other three blocks in the middle represent the rest of the network, based on the number of output channels of each layer: *block1* contains all the layers with an output of 16 channels, *block2* contains all the layers with an output of 32 channels and *block3* collects the layers with an output of 64 channels.

ResNet is an already optimized architecture and so it is more challenging to prune compared to, e.g. VGG. Nevertheless, SeReNe is still able to prune about 40% of the neurons and 70% of the parameters over the original ResNet-32. This is reflected on the size of the network, which drops from 1.84 MB (1.63 MB compressed) to 0.87 MB (0.57MB compressed).

G. AlexNet on CIFAR-100

Next, we up-scale in the output dimensionality of the learning problem, i.e. in the number of classes C , testing the proposed method on an AlexNet-like network over the CIFAR-100 dataset. Such dataset consists of 32×32 RGB images divided in 100 classes (50k training images, 10k test images). In this experiment we use SGD with learning rate $\eta = 0.1$ and $PWE = 20$ epochs. Concerning SeReNe (LB), we used $\lambda = 10^{-5}$ and $TWT = 1.5$ and the pruning process lasted 300 epochs.

Table VIII shows compression ratios in excess of 179x, whereas the network size drops from 92.31 MB to 43.80 MB and further to 2.47 MB after compression.

With respect to CIFAR-10, we hypothesize that the larger number of target classes to discriminate prevents pruning neurons in the convolutional layers, yet it allows to prune a significant number of neurons from the hidden fully connected layers. Contrarily from the previous experiments, the top-5 and the top-1 errors *improve* with respect to the baseline.

H. ResNet-101 on ImageNet

As a last experiment, we test SeReNe on ResNet-101 trained over ImageNet (ILSVRC-2012), using the pre-trained network

provided by the torchvision library.³

Due to the long training time, we employed a batch-wise heuristic such that, instead of waiting for a performance plateau, the pruning step is taken every time a fifth of the train set (around 7.9k iterations) has been processed. We trained the network using SGD with a learning rate $\eta = 0.001$ and momentum 0.9; for SeReNe (LB) we used $\lambda = 10^{-6}$ and $TWT = 0$.

Table IX shows the result of the pruning procedure with the layers grouped in blocks similarly as for the ResNet-32 experiment. Despite the complexity of the classification problem (1000 classes) that makes challenging pruning entire neurons, we prune around 86% of the parameters and obtain a network that is smaller in size, especially when compressed, going from 156.67 MB to only 27.84 MB. Comparing SeReNe (LB) to ℓ_2 -pruning we observe a boost in the performance when using SeReNe, but not as wide as in previous results. If the classification task to solve is complex (as for ImageNet), necessarily the sensitivity of many neurons will be high (because we need more neurons to solve the task). If many neurons will have a comparable sensitivity, then for those SeReNe's regularization naturally reduces to ℓ_2 regularization. This is one of the major benefits of SeReNe, which self-tunes the penalty to the neurons according to the models complexity and the complexity to solve the target task.

I. Experiments on mobile devices

As a last experiment, we benchmark some of the architectures pruned with SeReNe on an Huawei P20 smartphone equipped with 4x2.36 GHz Cortex-A73 + 4x1.84GHz Cortex-A53 processors and 4GB RAM, running Android 8.1 "Oreo". Table X shows the the inference time for ResNet-32, VGG-16 and AlexNet (all figures are obtained averaging 1,000 inferences on the device). SeReNe-pruned architectures show consistently lower inference time in the light of the fewer neurons in the pruned network, with a top speedup for VGG-16 in excess of a 2x factor. These results do not account for strategies commonly employed to boost inference speed, like parameters quantization or custom libraries for sparse tensors processing. We hypothesize that such strategies, being orthogonal to neuron pruning, would further boost inference time.

VI. CONCLUSIONS

In this work we have proposed a sensitivity-driven neural regularization technique. The effect of this regularizer is to penalize all the parameters belonging to a neuron whose output is not influential in the output of the network. We have learned that the evaluation of the sensitivity at the neuron level (SeReNe) is extremely important in order to promote a structured sparsity in the network, being able to obtain a smaller network with minimal performance loss. Our experiments show that the SeReNe strikes a favorable trade-off between ability to prune neurons and computational cost, while controlling the impairment in classification performance. For all the tested architectures and datasets, our

³<https://pytorch.org/docs/stable/torchvision/models.html>

TABLE VIII: AlexNet trained on CIFAR-100.

Approach	Remaining parameters (%) [neurons]								Compr. ratio	Network size [MB]		Top-1 (%)	Top-5 (%)
	Conv1	Conv2	Conv3	Conv4	Conv5	FC1	FC2	FC3		.onnx	.7z		
Baseline	100	100	100	100	100	100	100	100	1x	92.31	→ 79.27	45.58	20.09
	[64]	[192]	[384]	[256]	[256]	[4096]	[4096]	[100]					
Sparse VD [23]	-	-	-	-	-	-	-	-	26.45x	-		49.62	20.93
ℓ_2 +pruning	75.00	21.95	5.21	3.65	5.59	0.62	0.17	6.44	114.45x	60.88	→ 3.56	46.43	19.91
	[64]	[192]	[384]	[256]	[256]	[4094]	[2180]	[100]					
SeReNe (LB)	79.05	20.33	5.72	3.33	2.23	0.18	0.04	2.77	179.52x	43.80	→ 2.47	44.99	17.88
	[64]	[191]	[384]	[256]	[256]	[3322]	[1310]	[100]					

TABLE IX: ResNet-101 trained on ImageNet.

Approach	Remaining parameters (%) [neurons]							Compr. ratio	Network size [kB]		Top-1 (%)	Top-5 (%)
	Conv1	Block1	Block2	Block3	Block4	FC1	.onnx		.7z			
Baseline	100	100	100	100	100	100	1x	174.49	→ 156.67	22.63	6.44	
	[64]	[1408]	[3584]	[36352]	[11264]	[1000]						
Sparse VD [23]	-	-	-	-	-	-	2.48x	-		35.76	13.45	
ℓ_2 +pruning	53.12	25.42	25.57	13.71	17.74	51.94	5.75x	172.94	→ 32.93	28.33	9.18	
	[49]	[1241]	[3280]	[33278]	[11250]	[1000]						
SeReNe (LB)	55.36	24.27	23.79	11.24	14.81	40.82	6.94x	172.15	→ 27.84	28.41	9.45	
	[49]	[1197]	[3142]	[31948]	[11249]	[1000]						

TABLE X: Inference measures on Huawei P20.

Architecture	Approach	Inference time [ms]
ResNet-32	Baseline	32.12 ± 3.62
	SeReNe (LB)	24.83 ± 3.59
VGG-16 (VGG-1)	Baseline	204.21 ± 6.05
	SeReNe (LB)	98.67 ± 8.71
AlexNet	Baseline	131.41 ± 11.04
	SeReNe (LB)	75.27 ± 8.70

sensitivity-based approach proved to introduce a structured sparsity while achieving state-of-the-art compression ratios. Furthermore, the designed sparsifying algorithm, making use of cross-validation, guarantees minimal (or no) performance loss, which can be tuned by the user via an hyper-parameter (TWT). Future work includes deployment on physical embedded devices making use of deep network as well as using a quantization-based regularization jointly with the neuron sensitivity to further compress deep networks.

ACKNOWLEDGEMENT

This research was partially funded by Sisvel Technology, Turin, Italy. This work has received funding from the European Unions Horizon 2020 research and innovation programme under grant agreement No 825111, DeepHealth Project.

APPENDIX A J MADE EXPLICIT

In order to compute S given (5), we can proceed directly. However, this is problematic as it requires C different calls for the differentiation engine. Given (4), we can inspect whether we can reduce the computation by defining an overall objective

function which, when differentiated, yields S as a result. We name it J :

$$\begin{aligned}
 J &= \int S_{n,i} dp_{n,i} = \int \frac{1}{C} \sum_{k=1}^C \left| \frac{\partial y_{N,k}}{\partial p_{n,i}} \right| dp_{n,i} = \\
 &= \frac{1}{C} \int \sum_{k=1}^C \frac{\partial y_{N,k}}{\partial p_{n,i}} \text{sign} \left(\frac{\partial y_{N,k}}{\partial p_{n,i}} \right) dp_{n,i}. \quad (24)
 \end{aligned}$$

Here we can use Fubini-Tonelli's theorem:

$$\begin{aligned}
 J &= \frac{1}{C} \sum_{k=1}^C \int \frac{\partial y_{N,k}}{\partial p_{n,i}} \text{sign} \left(\frac{\partial y_{N,k}}{\partial p_{n,i}} \right) dp_{n,i} \\
 &= \frac{1}{C} \sum_{k=1}^C y_{N,k} \text{sign} \left(\frac{\partial y_{N,k}}{\partial p_{n,i}} \right). \quad (25)
 \end{aligned}$$

Unfortunately, we have no efficient way to compute $\text{sign} \left(\frac{\partial y_{N,k}}{\partial p_{n,i}} \right)$ and the only certain way is to compute $\frac{\partial y_{N,k}}{\partial p_{n,i}}$ directly, $\forall C$.

APPENDIX B

LENET300 WITH SIGMOID ACTIVATION ON MNIST

APPENDIX C

EXPLICIT DERIVATION FOR SERENE REGULARIZATION FUNCTION

Here we focus on the update rule (15): we aim at minimizing the overall objective function

$$O = \eta L + \lambda R, \quad (26)$$

where

$$R = \int w_{n,i,j} \bar{S}_{n,i} dw_{n,i,j}. \quad (27)$$

Here on we will drop the subscripts n, i, j . Let us consider the formulation of the sensitivity in (4): (27) becomes

$$R = \int w \cdot \bar{S} \cdot \Theta(\bar{S}) \cdot dw \quad (28)$$

where $\Theta(\cdot)$ is the one-step function and

$$\bar{S} = 1 - \frac{1}{C} \sum_{k=1}^C \left| \frac{\partial y_{N,k}}{\partial p} \right|. \quad (29)$$

We can re-write (28) as

$$\begin{aligned} R &= \int w \left[1 - \frac{1}{C} \sum_{k=1}^C \left| \frac{\partial y_{N,k}}{\partial p} \right| \right] \Theta(\bar{S}) dw \\ &= \int w \cdot \Theta(\bar{S}) \cdot dw - \int w \cdot \frac{1}{C} \sum_{k=1}^C \left| \frac{\partial y_{N,k}}{\partial p} \right| \cdot \Theta(\bar{S}) \cdot dw \\ &= \frac{w^2}{2} \cdot \Theta(\bar{S}) - \int w \cdot \frac{1}{C} \sum_{k=1}^C \left| \frac{\partial y_{N,k}}{\partial p} \right| \cdot \Theta(\bar{S}) \cdot dw. \end{aligned} \quad (30)$$

Let us define

$$J_R = - \int w \cdot \frac{1}{C} \sum_{k=1}^C \left| \frac{\partial y_{N,k}}{\partial p} \right| \cdot \Theta(\bar{S}) \cdot dw. \quad (31)$$

Considering that w is k -independent and that $\frac{1}{C}$ is a constant, we can write

$$J_R = - \frac{1}{C} \int \sum_{k=1}^C w \left| \frac{\partial y_{N,k}}{\partial p} \right| \cdot \Theta(\bar{S}) \cdot dw. \quad (32)$$

Here we are allowed to apply Fubini-Tonelli's theorem, swapping sum and integral:

$$\begin{aligned} J_R &= - \frac{1}{C} \sum_{k=1}^C \int w \left| \frac{\partial y_{N,k}}{\partial p} \right| \cdot \Theta(\bar{S}) \cdot dw \\ &= - \frac{1}{C} \sum_{k=1}^C \int w \frac{\partial y_{N,k}}{\partial p} \cdot \text{sign} \left(\frac{\partial y_{N,k}}{\partial p} \right) \cdot \Theta(\bar{S}) \cdot dw. \end{aligned} \quad (33)$$

Now integrating by parts:

$$\begin{aligned} J_R &= - \frac{1}{C} \sum_{k=1}^C \int w \frac{\partial y_{N,k}}{\partial p} \cdot \text{sign} \left(\frac{\partial y_{N,k}}{\partial p} \right) \cdot \Theta(\bar{S}) \cdot dw \\ &= - \frac{1}{C} \sum_{k=1}^C \left\{ \frac{w^2}{2} \frac{\partial y_{N,k}}{\partial p} \cdot \text{sign} \left(\frac{\partial y_{N,k}}{\partial p} \right) \cdot \Theta(\bar{S}) + \int \frac{w^2}{2} \cdot \frac{\partial}{\partial w} \frac{\partial y_{N,k}}{\partial p} \cdot \text{sign} \left(\frac{\partial y_{N,k}}{\partial p} \right) \cdot \Theta(\bar{S}) \cdot dw \right\}. \end{aligned} \quad (34)$$

According to the derivative chain rule, we can re-write (34) as

$$\begin{aligned} J_R &= - \frac{1}{C} \sum_{k=1}^C \left\{ \frac{w^2}{2} \frac{\partial y_{N,k}}{\partial p} \cdot \text{sign} \left(\frac{\partial y_{N,k}}{\partial p} \right) \cdot \Theta(\bar{S}) + \int \frac{w^2}{2} \cdot \frac{\partial^2 y_{N,k}}{\partial p^2} \cdot \frac{\partial p}{\partial w} \cdot \text{sign} \left(\frac{\partial y_{N,k}}{\partial p} \right) \cdot \Theta(\bar{S}) \cdot dw \right\}. \end{aligned} \quad (35)$$

Applying infinite steps of integration by parts we have in the end

$$\begin{aligned} J_R &= \frac{1}{C} \Theta(\bar{S}) \sum_{k=1}^C \text{sign} \left(\frac{\partial y_{N,k}}{\partial p} \right) \left[- \frac{w^2}{2} \frac{\partial y_{N,k}}{\partial p} + \sum_{i=1}^{\infty} (-1)^{i+1} \frac{w^{i+2}}{(i+2)!} \frac{\partial^{i+1} y_{N,k}}{\partial p^{i+1}} \frac{\partial^i p}{\partial w^i} \right]. \end{aligned} \quad (36)$$

Hence, the overall minimized R function is

$$\begin{aligned} R &= \Theta(\bar{S}) \left\{ \frac{w^2}{2} + \frac{1}{C} \sum_{k=1}^C \text{sign} \left(\frac{\partial y_{N,k}}{\partial p} \right) \left[- \frac{w^2}{2} \frac{\partial y_{N,k}}{\partial p} + \sum_{i=1}^{\infty} (-1)^{i+1} \frac{w^{i+2}}{(i+2)!} \frac{\partial^{i+1} y_{N,k}}{\partial p^{i+1}} \frac{\partial^i p}{\partial w^i} \right] \right\}. \end{aligned} \quad (37)$$

APPENDIX D DERIVATION OF (17)

Let us recall the formulation in (15). According to (14), we can re-write it as

$$w_{n,i,j}^{t+1} = w_{n,i,j}^t - \eta \frac{\partial L}{\partial w_{n,i,j}^t} - \lambda w_{n,i,j} (1 - S_{n,i}) \Theta(1 - S_{n,i}). \quad (38)$$

Given the definition of $S_{n,i}$ in (4), we can write

$$\begin{aligned} w_{n,i,j}^{t+1} &= w_{n,i,j}^t - \eta \frac{\partial L}{\partial w_{n,i,j}^t} + \\ &\quad - \lambda w_{n,i,j} \Theta(1 - S_{n,i}) \left(1 - \frac{1}{C} \sum_{k=1}^C \left| \frac{\partial y_{N,k}}{\partial p_{n,i}} \right| \right). \end{aligned} \quad (39)$$

We can multiply the insensitivity by the term $w_{n,i,j}$:

$$\begin{aligned} w_{n,i,j}^{t+1} &= w_{n,i,j}^t - \eta \frac{\partial L}{\partial w_{n,i,j}^t} + \\ &\quad - \lambda \Theta(1 - S_{n,i}) \left(w_{n,i,j} - \frac{1}{C} \sum_{k=1}^C \left| \frac{\partial y_{N,k}}{\partial p_{n,i}} \right| \cdot w_{n,i,j} \right). \end{aligned} \quad (40)$$

Finally here, observing (16), we find back (17).

APPENDIX E LENET300 WITH SIGMOID ACTIVATION ON MNIST

Finally, we repeat the experiment in Sec. V-B yet replacing the ReLU activations with sigmoids in the hidden layers. We optimize a pre-trained LeNet300 using SGD with learning rate $\eta = 0.1$, $PWE = 20$ epochs, $TWT = 0$ for target Top-1 errors of 1.7% (Table XI) and 1.95% (Table XII).

SeReNe achieves a sparser (and smaller) architecture than $\ell_2 +$ pruning for both error rates. Interestingly, for the 1.7% error rate, $\ell_2 +$ pruning is not able to prune any neuron, whereas SeReNe prunes 85 neurons from FC1, with a 10 times higher compression ratio. This reflects in the compressed model size: while $\ell_2 +$ pruning squeezes the architecture to 536kB, SeReNe compresses it to 75kB only.

TABLE XI: LeNet-300 trained on MNIST (sigmoid activation, 1.7% error rate).

Approach	Remaining parameters (%)			Compr. ratio	Remaining neurons	Model size [kB]			Top-1 (%)
	FC1	FC2	FC3			.onnx	.7z		
Baseline	100	100	100	1x	[300]-[100]-[10]	1043	→	963	1.72
ℓ_2 + pruning	46.27	82.87	97.90	1.97x	[300]-[100]-[10]	1043	→	536	1.75
SeReNe	2.44	16.73	85.30	22.31x	[215]-[100]-[10]	749	→	75	1.72

TABLE XII: LeNet-300 trained on MNIST (sigmoid activation, 1.95% error rate).

Approach	Remaining parameters (%)			Compr. ratio	Remaining neurons	Model size [kB]			Top-1 (%)
	FC1	FC2	FC3			.onnx	.7z		
Baseline	100	100	100	1x	[300]-[100]-[10]	1043	→	963	1.72
ℓ_2 + pruning	4.32	30.53	90.80	12.92x	[290]-[100]-[10]	1008	→	112	1.98
SeReNe	1.15	9.32	76.00	40.66x	[179]-[99]-[10]	624	→	45	1.95

REFERENCES

- [1] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *Advances in neural information processing systems*, 2012, pp. 1097–1105.
- [2] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *3rd International Conference on Learning Representations, ICLR 2015 - Conference Track Proceedings*, 2015.
- [3] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [4] T. M. P. E. Group, "Compression of neural networks for multimedia content description and analysis," MPEG 125 - Marrakesh.
- [5] Y. Lu, G. Lu, R. Lin, J. Li, and D. Zhang, "Srgc-nets: Sparse repeated group convolutional neural networks," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 31, pp. 2889–2902, 2020.
- [6] J. Wu, C. Leng, Y. Wang, Q. Hu, and J. Cheng, "Quantized convolutional neural networks for mobile devices," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 4820–4828.
- [7] L. Sagun, U. Evci, V. U. Guney, Y. Dauphin, and L. Bottou, "Empirical analysis of the hessian of over-parametrized neural networks," *6th International Conference on Learning Representations, ICLR 2018 - Workshop Track Proceedings*, 2018.
- [8] Y. LeCun, J. S. Denker, and S. A. Solla, "Optimal brain damage," in *Advances in neural information processing systems*, 1990, pp. 598–605.
- [9] D. P. Kingma, T. Salimans, and M. Welling, "Variational dropout and the local reparameterization trick," in *Advances in Neural Information Processing Systems*, 2015, pp. 2575–2583.
- [10] E. Tartaglione, S. Lepsøy, A. Fiandrotti, and G. Francini, "Learning sparse neural networks via sensitivity-driven regularization," in *Advances in Neural Information Processing Systems*, 2018, pp. 3878–3888.
- [11] S. Han, J. Pool, J. Tran, and W. Dally, "Learning both weights and connections for efficient neural network," in *Advances in neural information processing systems*, 2015, pp. 1135–1143.
- [12] J. Frankle and M. Carbin, "The lottery ticket hypothesis: Finding sparse, trainable neural networks," *7th International Conference on Learning Representations, ICLR 2019*, 2019.
- [13] M. Naumov, L. Chien, P. Vandermersch, and U. Kapasi, "Cusp sparse library," in *GPU Technology Conference*, 2010.
- [14] J. Bai, F. Lu, K. Zhang *et al.*, "Onnx: Open neural network exchange," <https://github.com/onnx/onnx>, 2019.
- [15] V. Lebedev and V. Lempitsky, "Fast convnets using group-wise brain damage," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 2554–2564.
- [16] B. Liu, M. Wang, H. Foroosh, M. Tappen, and M. Pensky, "Sparse convolutional neural networks," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2015, pp. 806–814.
- [17] M. Zhu and S. Gupta, "To prune, or not to prune: exploring the efficacy of pruning for model compression," *6th International Conference on Learning Representations, ICLR 2018 - Workshop Track Proceedings*, 2018.
- [18] W. Wen, C. Wu, Y. Wang, Y. Chen, and H. Li, "Learning structured sparsity in deep neural networks," in *Advances in neural information processing systems*, 2016, pp. 2074–2082.
- [19] T. Gale, E. Elsen, and S. Hooker, "The state of sparsity in deep neural networks," *CoRR*, vol. abs/1902.09574, 2019. [Online]. Available: <http://arxiv.org/abs/1902.09574>
- [20] M. Gong, J. Liu, H. Li, Q. Cai, and L. Su, "A multiobjective sparse feature learning model for deep neural networks," *IEEE transactions on neural networks and learning systems*, vol. 26, no. 12, pp. 3263–3277, 2015.
- [21] M. Lin, R. Ji, Y. Zhang, B. Zhang, Y. Wu, and Y. Tian, "Channel pruning via automatic structure search," in *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI-20*, C. Bessiere, Ed. International Joint Conferences on Artificial Intelligence Organization, 7 2020, pp. 673–679, main track.
- [22] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, "Dropout: a simple way to prevent neural networks from overfitting," *The Journal of Machine Learning Research*, vol. 15, no. 1, pp. 1929–1958, 2014.
- [23] D. Molchanov, A. Ashukha, and D. Vetrov, "Variational dropout sparsifies deep neural networks," in *Proceedings of the 34th International Conference on Machine Learning-Volume 70*. JMLR. org, 2017, pp. 2498–2507.
- [24] A. N. Gomez, I. Zhang, K. Swersky, Y. Gal, and G. E. Hinton, "Learning sparse networks using targeted dropout," *CoRR*, vol. abs/1905.13678, 2019.
- [25] G. Hinton, O. Vinyals, and J. Dean, "Distilling the knowledge in a neural network," *arXiv preprint arXiv:1503.02531*, 2015.
- [26] N. Papernot, P. McDaniel, X. Wu, S. Jha, and A. Swami, "Distillation as a defense to adversarial perturbations against deep neural networks," in *2016 IEEE symposium on security and privacy (SP)*. IEEE, 2016, pp. 582–597.
- [27] T. Li, J. Li, Z. Liu, and C. Zhang, "Few sample knowledge distillation for efficient network compression," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 14 639–14 647.
- [28] Y. Wang, X. Zhang, L. Xie, J. Zhou, H. Su, B. Zhang, and X. Hu, "Pruning from scratch," in *AAAI*, 2020, pp. 12 273–12 280.
- [29] N. Lee, T. Ajanthan, and P. Torr, "Snip: Single-shot network pruning based on connection sensitivity," *7th International Conference on Learning Representations, ICLR 2019*, 2019.
- [30] E. Tartaglione, A. Bragagnolo, and M. Grangotto, "Pruning artificial neural networks: A way to find well-generalizing, high-entropy sharp minima," in *Artificial Neural Networks and Machine Learning – ICANN 2020*. Cham: Springer International Publishing, 2020, pp. 67–78.
- [31] C. Louizos, M. Welling, and D. P. Kingma, "Learning sparse neural networks through L_0 regularization," *6th International Conference on Learning Representations, ICLR 2018 - Conference Track Proceedings*, 2018.
- [32] E. Tartaglione, D. Perlo, and M. Grangotto, "Post-synaptic potential regularization has potential," in *International Conference on Artificial Neural Networks*. Springer, 2019, pp. 187–200.
- [33] K. Ullrich, M. Welling, and E. Meeds, "Soft weight-sharing for neural network compression," *5th International Conference on Learning Representations, ICLR 2017 - Conference Track Proceedings*, 2019.
- [34] Y. Guo, A. Yao, and Y. Chen, "Dynamic network surgery for efficient dnns," *Advances in Neural Information Processing Systems*, pp. 1387–1395, 2016.

- [35] I. Pavlov, "Lzma sdk (software development kit)," 2007. [Online]. Available: <https://www.7-zip.org/sdk.html>
- [36] Y. LeCun, L. Bottou, Y. Bengio, P. Haffner *et al.*, "Gradient-based learning applied to document recognition," *Proceedings of the IEEE*, vol. 86, no. 11, pp. 2278–2324, 1998.
- [37] H. Xiao, K. Rasul, and R. Vollgraf, "Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms," *CoRR*, vol. abs/1708.07747, 2017. [Online]. Available: <http://arxiv.org/abs/1708.07747>



Enzo Tartaglione (M'20) received the joint MS degree in Electronic Engineering from Politecnico di Torino, University of Illinois at Chicago and Politecnico di Milano in 2015, with 110/110 cum laude. In 2019 he received the Ph.D. degree from Politecnico di Torino, cum laude, with the thesis "From Statistical Physics to Algorithms in Deep Neural Systems". He is currently a postdoc at Università degli Studi di Torino. His main research interests are neural network applied to medical image processing, unbiased learning and compression/pruning.



Andrea Bragagnolo received his Master's Degree in Computer Science from the University of Turin, Italy, in 2019. He is currently a Ph.D. student at the University of Turin and employed by Synesthesia for the development of smart systems for retail workers. His current research interests include pruning and simplification of Neural Networks and synthetic datasets for Object Detection.



skills in the industry.

Francesco Odierna is a Computer Scientist with a focus on Machine Learning and Deep Learning technologies. He holds a master's degree in Computer Science from the University of Turin with a thesis on Neural Networks Sparsification and a bachelor's degree in Computer Science from the University of Salerno. While obtaining his master's degree, he also worked at Télécom ParisTech, a challenge in which he based his thesis research. He enrolled in a Ph.D. in Deep Learning at the University of Pisa but decided to put it on hold to further develop his



Attilio Fiandrotti (M'12) received his Ph.D. in computer science from Politecnico di Torino in 2010. He is currently assistant professor at Università di Torino, Dipartimento di Informatica, Italy, and holds a position as *matre de conférences* with the Multimedia Group at LTCI, Télécom Paris, Institut Polytechnique de Paris. His current research interests focus mainly on deep learning techniques for image and video analysis, compression, and synthesis and the distribution of multimedia contents over packet networks.



Marco Grangetto (S'99-M'03-SM'09) received the Electrical Engineering degree and the Ph.D. degree from the Politecnico di Torino, Italy, in 1999 and 2003, respectively. He is currently a Full Professor with the Computer Science Department, Università di Torino, where he coordinates research in the area of image processing and computer vision. His research interests are in the fields of multimedia signal processing and networking. His expertise includes wavelets, image and video coding, data compression, video error concealment, error resilient video coding, computer vision and biomedical image processing. He was the recipient of the Premio Optime by Unione Industriale di Torino in September 2000 and a Fulbright Grant in 2001 for a research period with the Department of Electrical and Computer Engineering, University of California, San Diego, USA. He participated in the ISO standardization activities on Part 11 of the JPEG2000 standard. He was associate editor of IEEE Transactions on Communications, currently he is associate editor of IEEE Transactions on Multimedia. He is a member of Moving Picture, Audio and Data Coding by Artificial Intelligence (MPAI).