



Structured Ensembles: An approach to reduce the memory footprint of ensemble methods

Jary Pomponi ^{*}, Simone Scardapane, Aurelio Uncini

Department of Information Engineering, Electronics and Telecommunications (DIET), Sapienza University of Rome, Italy

ARTICLE INFO

Article history:

Received 25 May 2021

Received in revised form 24 July 2021

Accepted 3 September 2021

Available online 16 September 2021

Keywords:

Ensemble

Continual learning

Pruning

Structured pruning

Neural networks

Deep learning

ABSTRACT

In this paper, we propose a novel ensembling technique for deep neural networks, which is able to drastically reduce the required memory compared to alternative approaches. In particular, we propose to extract multiple sub-networks from a single, untrained neural network by solving an end-to-end optimization task combining differentiable scaling over the original architecture, with multiple regularization terms favouring the diversity of the ensemble. Since our proposal aims to detect and extract sub-structures, we call it **Structured Ensemble**. On a large experimental evaluation, we show that our method can achieve higher or comparable accuracy to competing methods while requiring significantly less storage. In addition, we evaluate our ensembles in terms of predictive calibration and uncertainty, showing they compare favourably with the state-of-the-art. Finally, we draw a link with the continual learning literature, and we propose a modification of our framework to handle continuous streams of tasks with a sub-linear memory cost. We compare with a number of alternative strategies to mitigate catastrophic forgetting, highlighting advantages in terms of average accuracy and memory.

© 2021 Elsevier Ltd. All rights reserved.

1. Introduction

An ensemble is the combination of multiple models, trained to solve the same task and joined to improve the overall performance. Informally, in a good ensemble each member is accurate on its own, but makes independent errors during the prediction (Perrone & Cooper, 1992). By combining the outputs of its components, an ensemble can achieve better performance than any base member. For neural networks, a simple strategy for building an ensemble of models is to train repeatedly the same architecture starting from different initial conditions, and averaging the resulting predictions. In this setup, diversity is provided by the exponentially growing number of possible minima (Kawaguchi, 2016), while averaging provides a degree of smoothing against minima with poor generalization capability.

While this *naive* strategy has been shown to provide significant advantages (e.g., Gustafsson, Danelljan, and Schon (2020), Krogh and Vedelsby (1995), Lakshminarayanan, Pritzel, and Blundell (2017) and Ovadia et al. (2019)), it is not widely used in practice. The computational time required to train multiple networks and the memory cost needed to save them limit its use. In fact, each model must be trained independently and stored in the system, and this can be infeasible when a single model

is composed by millions of parameters, or when the hardware capacity is limited (e.g., mobile devices).

With the emergence of highly over-parameterized models, a number of alternative techniques, e.g., Huang et al. (2017) and Wen, Tran, and Ba (2020), have been proposed to mitigate the space and time constraints. Key to many proposals is that, for a sufficiently large model, different components of the ensemble can share portions of the original network, opening the possibility of a sub-linear scaling of the memory with respect to the size of the ensemble (Wen et al., 2020). We briefly review these proposals in Section 2.1. In general, one can observe a trade-off between the amount of sharing (i.e., the memory cost) and the performance of the ensemble. Whether this gap can be closed remains an open research question motivating this paper.

Ensembles of deep neural networks (called *deep ensembles* in the following) are of particular interest in the field of continual learning (CL). In CL, a single network must solve multiple tasks sequentially, without forgetting already learned information (Thrun, 1998). This property is hard to achieve and yet crucial for real-world scenarios, where the domain can change. When this is not possible and some, or all, the past learned information is lost, we have a phenomenon called catastrophic forgetting (CF, French (1999) and McCloskey and Cohen (1989)). Because of its importance, many algorithms were proposed recently to achieve effective CL with neural models. Among these techniques (briefly reviewed in Section 2.2), several solutions find sub-structures within the original network, each one capable of

^{*} Corresponding author.

E-mail address: jary.pomponi@uniroma1.it (J. Pomponi).

solving a single task while leaving space for future tasks; usually, this is done by masking a portion of the network that is not used. In this sense, these methods can be seen as an ensemble of sub-networks, each one specialized on a single task. Thanks to this neat separation, many of these methods are capable of mitigating CF (Parisi, Kemker, Part, Kanan, & Wermter, 2019).

Deep ensembles and CL are problems that need to be addressed in order to achieve stable and robust agents that can operate autonomously in real-world applications and, as described before, methods developed in the two fields share a number of similarities. In this paper, we first develop in Section 3 a framework that is used for reducing the memory footprint of deep ensembles, without sacrificing the performances. Differently from other proposals in the literature, we optimize the members of the ensemble at *initialization*, by solving an optimization problem defined over partial sub-masks of the original network. Then, we extend the method in Section 3.3 for CL scenarios. We describe our contributions in more detail next.

1.1. Contributions of the paper

1.1.1. Structured ensembles

To build an ensemble of N networks, we apply N scaling vectors over the output of each inner and trainable layer in the original model. Each scaling vector is associated to a component of the ensemble. These vectors are trained by considering the network at initialization, while the networks' weights remain unchanged. In the end, a scaling value is used to estimate how important the associated neuron is for the given component of the ensemble. Only the most important ones are collected and the sub-network extracted, ignoring the others. This method is well motivated by the fact that, usually, a network is over-parameterized, and contains multiple sub-networks that can be extracted by its structure (Ramanujan, Wortsman, Kembhavi, Farhadi, & Rastegari, 2020), as demonstrated by the recent Lottery Tickets Hypothesis (Frankle & Carbin, 2018) and by several structured pruning strategies (He, Zhang, & Sun, 2017).

The overall procedure, which can be seen as a discovery of hidden paths in the original neural network, is the following: (1) the scaling vectors are applied and trained for a few epochs, (2) for each component $i = 1 \dots N$ we collect the scaling vector at position i in each layer of the original network, (3) the neurons which are important for the current sub-network are kept, and the others ignored. The process produces N smaller networks, and the dimension of each one can be regulated by choosing the percentage of how many neurons must be discarded in the extraction phase. An important property is that each network produced using this process is capable of achieving good performance despite the reduced number of parameters (even with a large percentage of discarded neurons), and when combining all the extracted networks into an ensemble the achieved score overcomes the one obtained using only the original network, due to the diversity of the networks that compose the ensemble. In order to create sub-networks that are all different from each other, we force a level of dissimilarity between scaling vectors during training using a maximum mean discrepancy (MMD) regularization term (Pomponi, Scardapane, & Uncini, 2021). To the best of our knowledge, the proposed method is the first which is capable of extracting a set of different sub-networks from a single untrained one.

We evaluate the efficiency of the proposed *Structured Ensemble* algorithm over different choices of architectures and datasets, and by investigating its robustness under different percentages of kept units. Moreover, we show that the ensembles created using Structured Ensemble can be better calibrated than other approaches and the performances remain good when it comes to evaluating the network on out-of-distribution samples. We also perform a number of ablation studies on the algorithm, including over the choice of initialization strategy for the scaling vectors.

1.1.2. Continual learning

As the second main contribution of the work, we extend Structured Ensemble to devise a novel approach for mitigating CF. In the CL scenario, the key idea is to associate each member of the ensemble (i.e., using a binary mask m_i) to a single task, by finding which portion of the network is maximally responsible for the performance on that task. In this way, we can remove CF by freezing these weights during the training of future tasks.

The overall procedure is the following: (1) before solving a task, we identify a smaller sub-network that can be used to achieve good performance, but without extracting it, (2) during the training only the weights in this subset are modified, while the others are masked to zero, or used but not changed if they are part of a sub-network associated to a past task. In this way, when we want to classify a sample from a given task, only the weights up to that task are used, and the others ignored. In the experiments, we evaluate empirically how Structured Ensemble performs on different CL scenarios, showing that the final accuracy achieved outperforms the other methods and that the additional memory required to store the binary masks is negligible; in the end, we show also how the hyper-parameters affect the final results, by performing different ablation studies.

2. Related work

In this section, we provide a short background on related works concerning ensembling 2.1, continual learning 2.2, and structured pruning 2.3.

2.1. Ensembles

Deep ensembles (Lakshminarayanan et al., 2017), originally inspired by bagging (Breiman, 1996), are models which combine multiple neural networks trained from randomly initialized weights. Ensembles have been theoretically studied to understand how and why the performances improve, as well as how to reduce the memory footprint and how to speed up the training time, often focusing on one of the two aspects. Regarding the memory aspect, the networks can be distilled into smaller ones, without loss of performance (Hinton, Vinyals, & Dean, 2015), or compressed, as exposed in Bucilua, Caruana, and Niculescu-Mizil (2006). Regarding the training time aspect, instead, recently in Huang et al. (2017) the authors proposed Snapshot Ensemble, in which a single model is forced to explore multiple local minima using cyclic learning rate schedules; the models obtained in these minima are saved and combined to build the final ensemble. Also, to reduce the training time, in Garipov, Izmailov, Podoprikin, Vetrov, and Wilson (2018) the geometric properties of a loss function are used to connect different minima and extract multiple models, used to build the final ensemble.

In addition to these methods, some techniques to improve at the same time memory footprint and training time have been proposed. Recently, in Wen et al. (2020) the authors proposed Batch Ensemble, a method which uses the Hadamard product to change the weights and the output of each layer using multiple sets of trainable scaling vectors, each one associated to a different ensemble's component. Similarly, in Lee, Purushwalkam, Cogswell, Crandall, and Batra (2015), the authors proposed a family of tree structured networks, called TreeNet, that simulates an ensemble by building a network consisting of zero or more shared initial layers, followed by a branching point and zero or more independent layers. Another approach is called MC-Dropout, which uses dropout to simulate multiple networks, by dropping the weights also during the forward step in the inference phase (Gal & Ghahramani, 2016). Both approaches are relaxations of the original ensemble formulation, and there is

usually a deterioration of the performances (if compared to the original deep ensemble models).

Recently, some studies provide more insight into the deep ensemble approach, studying why and how this method performs better than others. In Allen-Zhu and Li (2020) the authors found that ensemble in deep learning works very differently from traditional learning theories, and, to understand and explain the hidden mechanisms of deep ensembles, they develop a new theory, while in Fort, Hu, and Lakshminarayanan (2019) the ensemble approach is studied from a loss landscape perspective by measuring the similarity between the different models which compose the ensemble.

Other important aspects, that justify studying the ensemble approaches, are that these models are usually more calibrated than a single neural network (Guo, Pleiss, Sun, & Weinberger, 2017) and that it is possible to calculate a measure of uncertainty associated to a prediction. Both aspects are crucial in real-world applications, e.g., driver-less vehicles and medical applications (Kwon, Won, Kim, & Paik, 2020).

2.2. Continual learning

A continual learning (CL) scenario is one in which a model is trained to solve a number of sequential tasks, and should not forget how to solve past tasks while learning how to solve the current one (Thrun, 1998). When the models lose the ability to solve past tasks we have a phenomenon called catastrophic forgetting (CF (French, 1999; McCloskey & Cohen, 1989)). A CL method should be capable of alleviating, or removing, CF while solving efficiently the current task.

A CL scenario is a set of tasks and rules on how these tasks are acquired and which information can be retrieved from them. Many scenarios have been formulated, and we follow those proposed in Van de Ven and Tolias (2019). More specifically, in this paper we focus on the scenario called Task-IL, where a task's identity is given during both training and testing, and the tasks contain disjoint classes. A generic model used to solve this kind of scenario is composed of a backbone network, that is shared between all tasks and therefore may suffer from CF, and a set of smaller networks, called *solvers*, one per task, that are used to classify only the samples from the associated task, that are trained on top of the backbone. Despite its apparent simplicity, many approaches have been developed to solve this scenario.

Following the categorization proposed in Parisi et al. (2019), the methods to mitigate or remove CF in this scenario can be divided into three broad categories: **Rehearsal strategies**, **Regularization techniques**, and **Architectural strategies** (our proposal falls among the architectural category, on which we focus more). The **Architectural strategies** group contains all the methods that operate over the architecture of a neural network (e.g. custom layers, weights freezing/pruning) to mitigate or remove CF. Even if this group of methods has received less attention if compared to the others, many interesting approaches exist. In Rusu et al. (2016) the authors proposed Progressive Neural Networks, a model which is capable of expanding its size when a new task is encountered. It removes CF, but the required memory can be prohibitive. Batch Ensembles (Wen et al., 2020), mentioned before, computes a set of scaling vectors associated to the different tasks, training each of them independently. This approach requires small additional memory, but it is limited because the backbone is trained only during the first task, while for the other tasks the backbone is fixed and only the additional vectors are trained, limiting the final performance. Another interesting approach, proposed in Golkar, Kagan, and Cho (2019), is based on the classic pruning scheme, in which the most important weights for solving a task are saved and freed, while the others

are used to solve new tasks. The main problems of this method are that it is not possible to automatically decide the pruning percentage, and some accuracy can be lost when pruning the final model; also, past information is not used, limiting the achieved performance. The closest method to the one we propose here is Supermasks In Superposition (Wortsman et al., 2020), where for each task a binary mask over the backbone is learned. Our method has some similarities with these methods, but it is deeply different as a whole. In fact, it can be seen as a structured pruning approach, in which the masks are retrieved by searching for the best structure, instead of pruning the redundant weights after the training, and each structure is well separated from the others but reused to extract useful information from future samples, without compromising the results on past tasks.

2.3. Structured pruning

Our proposal has a strong connection with structured pruning approaches (He et al., 2017; Li, Kadav, Durdanovic, Samet, & Graf, 2016). These methods do not prune each weight singularly, but entire groups of them together (e.g., neurons, filters, or channels). Usually, these methods add a regularization term to the optimization process. Then, when the minimization is over, the groups are pruned. Our proposal, instead, is capable of pruning the neurons before the training process. To this end, we extend the SNIP approach from Lee, Ajanthan, and Torr (2018), proposing a method that is capable of detecting multiple structured patterns within a neural network, used together to build up an ensemble of models, that is capable of solving efficiently a classification problem, without loss of performances.

3. Proposed method: structured ensemble

As described before, deep ensembles of multiple neural networks suffer from an expensive memory footprint. In this section we introduce our method, which aims to reduce the required memory by creating smaller networks from an untrained one. Firstly, we describe how the main idea works for reducing the dimension of a network. Secondly, we extend the method so that it can be used to extract multiple smaller networks within a single pre-processing procedure.

3.1. Extracting a sub-network

Let us consider a neural network which is composed by L layers L_l with $l = 1 \dots L$. To extract a substructure we want to consider only the layers that modify the dimension of its input, such as linear or convolutional layers. For example, let L_l be a generic linear layer with input size i and output o (the same can be done with a convolutional layer), that has weights matrix $W^l \in \mathbb{R}^{i \times o}$, and may have bias vector $b^l \in \mathbb{R}^o$, otherwise $b^l = 0$. Given an input $x \in \mathbb{R}^i$, we have the corresponding output $z^l = L_l(x) = W^{lT}x + b^l$.

Our intuition is that a matrix of weights contains a subset of the output neurons that can be extracted, and used to create a new layer with reduced dimensions. To estimate the importance of an output neuron in a layer L_l , we first multiply the output z^l by a learnable real valued vector, called scaling vector, $s^l \in \mathbb{R}^o$, so that the new output is given by $\bar{z}^l = L_l(x) \odot s^l$. We initialize s^l from the normal distribution $\mathcal{N}(0, 1)$.

Denote by $S = \bigcup s^l$ the union of all scaling vectors through the network, and by $f(x; S)$ the output of the original network where every layer is replaced by its scaled version \bar{z}^l . We can find S by solving an optimization problem over the scaling vectors:

$$\min_S \mathbb{E}[\mathcal{L}(y, f(x; S))] \quad (1)$$

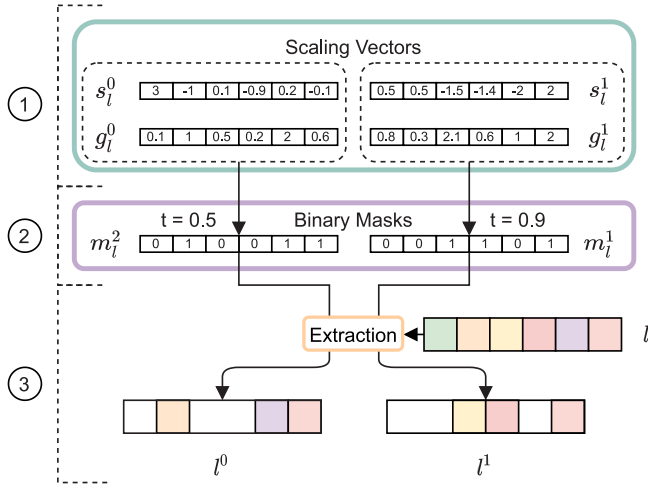


Fig. 1. The image shows the extraction process of two linear layers from a single one l , using pruning percentage $p = 0.5$. For simplicity, instead of showing the whole matrix of weights, only the output neurons of the layer are shown. The process starts in the step 1: after the minimization step, the gradients associated to the scaling vectors are collected. In the step 2 the gradients are used to calculate the thresholds th , one for each scaling vector, and thus the sets of indexes are created. These sets are used in the last step 3, to create two new layers l^0 and l^1 , by extracting the neurons in the original layer l associated to indexes in the sets.

where \mathcal{L} is the loss function for our problem, e.g., cross-entropy. We found that (1) requires only a small number of epochs to train the scaling vectors to a reasonable value. Thus, it can be trained significantly faster than the original optimization problem over the weights of the network, which is necessary for keeping the approach bounded in terms of computational time.

Once we have the set S , the scaling vectors are used to measure how much each output neuron is important for solving the original problem. Directly using the final scaling values in a scaling vector s^l is not a good choice, since it does not take into consideration the correlation between the weights of a layer and its scaling values (higher scaling values do not imply more importance). To do this we need another estimator of this importance, and for Structured Ensemble we use the one proposed in Lee et al. (2018). Given an input sample x and its associated label y , the loss function $\mathcal{L}(\cdot, \cdot)$ and a scaling vector s^l associated to a linear layer $L_l(x)$, we calculate the gradient associated to a scaling value j as:

$$g_j^l = g(s_j^l) = \frac{\left| \frac{\partial \mathcal{L}(y, f(x; S))}{\partial s_j^l} \right|}{\sum_k^L \sum_i^N \left| \frac{\partial \mathcal{L}(y, f(x; S))}{\partial s_i^k} \right|} \quad (2)$$

The gradient value $g(\cdot)$ can be computed using the loss over a single sample, a batch of samples, or the averaged gradients computed over the whole dataset. Differently from Lee et al. (2018), in Structured Ensemble we use the value $g(\cdot)$ on the whole dataset, to avoid sub-optimal extraction of the most important neurons.

Once all the values $g(\cdot)$ are computed, we collect them into a set $G = \bigcup_l \{g(s_i^l) | s_i^l \in s^l\}$. To extract a sub-structure we need to define a threshold over the values in G : if a value $g(s_i^l)$ is higher than the threshold we want to keep the neuron i of the layer l also in the extracted network, otherwise it can be discarded. In Structured Ensemble, the user sets the percentage of neurons p to discard (as a hyper-parameter), and the threshold is automatically calculated over the whole set of values as $th = Q_G(p)$. The function $Q_A(p)$ returns the percentile associated to the $p\%$ value of the set A . Once the threshold is calculated, we need to iterate over the

Algorithm 1 Structured Ensemble (pseudo-code).

Input untrained network $y = f(x)$, a dataset $\mathcal{D} = \{(x, y)\}$, the number of subnetworks to extract N , and a pruning percentage p .

Output a set of networks E .

- 1: Initialize the masks S_i , for $i = 1 \dots N$ and create $S = \bigcup S_i$.
- 2: Solve the minimization problem:

$$\min_S \mathbb{E}[\mathcal{L}(y, f(x; S))] + \frac{2\lambda}{(N(N-1))} \sum_{i=1, j>i}^N R(S_i, S_j)^{-1}.$$

- 3: $E = \{\}$
- 4: **for** $i = 1 \dots N$ **do**
- 5: Retrieve S_i
- 6: Given the dataset D , calculate the gradients g^l for each mask in S_i (Eq. (2)).
- 7: Calculate the threshold $th = Q_G(p)$ and l_i for each layer l (Eq. (3)).
- 8: Extract the sub-network $\tilde{f}(\cdot)$ of $f(\cdot)$, composed by the neurons that have $l_{i,l} = 1$.
- 9: $E \leftarrow E \cup \{\tilde{f}(\cdot)\}$.
- 10: **end for**
- 11: Return the ensemble of networks E .

layers of the network and extract the most important neurons. To do this, define the set of indexes associated to important neurons in the layer:

$$l_i = \{i | g_i^l > th\} \quad i = 1 \dots o \quad (3)$$

for each layer in the network l , where o is its output size. Then, if the layer l is the first one ($l = 1$), the new layer \bar{L}_l has a new weights matrix that is defined as $\bar{W}^l \in \mathbb{R}^{i \times |l_i|}$, where the following equivalence holds $\bar{W}_{j,k}^l = W_{j,k}^l$ for each $j = 1 \dots i$ and for each $k \in l_i$. The equivalence means not only that the weights have not been changed during the training of the scaling vectors, but also that the input size of the layer is unchanged, while the output's size is composed only of the neurons that have been selected to be part of the new layer; the same process must be applied to the bias vector, if present. The extraction procedure for a linear layer is visualized in Fig. 1.

When the layer is not the first one, in addition to the exposed procedure, we extract also the input neurons associated to the selected output's neurons of the previous layer using the set l_{l-1} , further reducing the dimension of the layer l . The extraction process is the same used before, but applied to the input ones.

The process is iterated for all layers whose dimensions depend on the precedent, or subsequent, layer. In this way, changing a layer may have repercussions also on connected layers, further reducing the number of parameters in the extracted neural network.

To conclude, the set of new layers \bar{L}_l compose the new neural network $f_0(x)$, which has the same number of layers and connections between them, but fewer weights.

3.2. Extracting multiple sub-networks via diversity regularization

The process described up to this point works for extracting one sub-network, but we want to extract N sufficiently diverse sub-networks to build up an ensemble model. To this end, we use a regularization term $R(\cdot, \cdot)$ to promote the distance between different scaling vectors in a layer, and thus the extracted networks.

Instead of applying a single scaling vector s^l on a generic layer l , multiple vectors $s^{l,i}$ are applied, with $i = 1 \dots N$, where N is

the number of networks to be extracted. To speed up the pre-processing, we train all the scaling vectors in parallel, following the idea proposed in Wen et al. (2020). Given a mini-batch of B samples, we partition it into B/N subsets, and apply a different set of scaling vectors to each subset. If B contains fewer samples than N or its dimension cannot be divided in N equally sized sets, we replicate some input samples to make up the difference, in such a way that a scaling vector does not see a sample twice in the same optimization step.

Denoting by $S_i = \bigcup s^{l,i}$ the set of scaling vectors associated to the i th component, and by $S = \bigcup S_i$ the set of all the scaling vectors, we can find N sub-networks in parallel by solving:

$$\min_S \mathbb{E}[\mathcal{L}(y, f(x; S))] + \frac{2\lambda}{(N(N-1))} \sum_{i=1, j>i}^N R(S_i, S_j)^{-1} \quad (4)$$

where $R(\cdot, \cdot)$ is a regularization term forcing a desired level of diversity between a pair of masks, λ is a hyper-parameter balancing the two terms, and $f(x; S_i)$ is the appropriately scaled version of f . In particular, we regularize the training by applying the MMD distance (Gretton, Borgwardt, Rasch, Schölkopf, & Smola, 2012) between pairs of vectors in a layer. Since the MMD distance is symmetric, the comparisons are reduced to $\frac{N(N-1)}{2}$ for each layer. The resulting distances are summed, weighted, and added to the loss which is minimized during the training of the scaling vectors. Formally, given a layer l and two scaling vectors $s^{l,i}$ and $s^{l,j}$, we calculate the distance as:

$$\begin{aligned} \text{MMD}_k^2(s^{l,i}, s^{l,j}) &= \frac{1}{n(n-1)} \sum_{z \neq k} \kappa(s_z^{l,i}, s_k^{l,i}) \\ &+ \frac{1}{n(n-1)} \sum_{z \neq k} \kappa(s_z^{l,j}, s_k^{l,j}) \\ &- \frac{2}{n^2} \sum_{k,z} \kappa(s_z^{l,i}, s_k^{l,j}) \end{aligned}$$

where n is the dimension of the scale vector and κ a kernel function. In the experiments we use the RBF kernel $\kappa(a, b) = \exp\{-\frac{\|a-b\|^2}{n}\}$.

After the training of the scaling vectors, extracting the sub-networks proceeds as exposed before (Section 3.1), but it is done N times: for each $i = 1 \dots N$, we use the set S_i to calculate the neurons to keep, we extract the associated sub-network $f_i(x)$ and we save it in the ensemble model. During this process the original network is not changed, because this would compromise future extractions. Fig. 2 shows the extraction process of two new layers from a generic linear layer l .

The final ensemble is the set $E = \{f_i | i = 1 \dots N\}$, whose predictions can be combined in standard ways (e.g., the average for regression and a majority vote for classification). The complete procedure to extract multiple subnetworks from an untrained one, while regularizing the distance between the masks, is summarized in Algorithm 1.

3.3. Structured ensemble as an approach to continual learning

As said before, we use the Task-IL CL scenario, proposed in Lopez-Paz and Ranzato (2017). In this scenario we receive a sequence of disjoint tasks $t = 1, \dots, M$, where a new task can be collected only when the current one is over. Each task is composed by a set of triples $\{(x_i, y_i, t)\}_{i=1}^S \in \mathcal{X} \times \mathbb{N}^+ \times \mathcal{Y}$, where x_i is a sample from the task (which is not present in other tasks), its label is y_i , which is defined with respect to the current task t , and t is an index identifying the current task. We

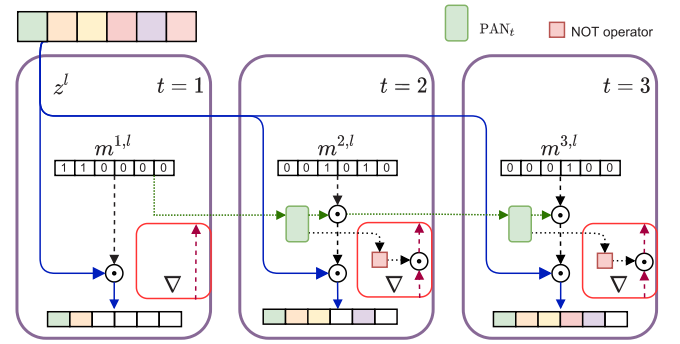


Fig. 2. The image shows how the masks are combined to mask the output z^l produced by a generic linear layer l . We follow the process for three tasks t , showing which output neurons and gradients are masked to avoid the changing of weights important to past tasks. In the first one the output vector is masked only using $m^{1,l}$, and the gradients ∇ are unmodified. In the following tasks, past masks are used to build the PAN_t set, that is combined with $m^{t,l}$ to mask the output of the layer, while the gradients are masked using the negation of the masks in PAN_t .

also define the neural network function $f(x; w, w_t)$, that given a sample x produces the associated output y . The backbone is shared between all the tasks and its weights set is w , while the classifier head used to calculate the prediction depends on the task t to which x belongs, and its weights set is w_t . This separation is crucial, since in this scenario CL may arise only in the backbone network, and thus the weights w_t in the solvers are not regularized by the methods.

Our main intuition is the following: neurons trained on past tasks are capable of extracting information useful also for future tasks, and this can be exploited to use a smaller trainable structure, instead of the full network. In this way, a network is partitioned into logical smaller networks using binary masks, keeping weights used to solve past tasks unchanged (so that no CF occur), while solving new tasks. As in the ensemble approach, when we talk about neurons we refer to the output neurons of each inner layer of the backbone (linear or convolutional), in this way the binary masks are not applied on the weights, but only on the layers' outputs, further reducing the required memory.

When a new task t is retrieved, we follow the same process for extracting a single network exposed before (Section 3.1), but without extracting it. Instead, we define the set of Active Neurons of t as $AN_t = \bigcup_{i=1}^t m^{i,l}$, where each binary mask $m^{t,l}$ is defined over the corresponding layer of the backbone as:

$$m_i^{t,l} = \begin{cases} 1 & i \in I_t \\ 0 & i \notin I_t \end{cases}$$

where I_t is calculated as in (3), using samples from task t , and the dimension of the mask is $m^t \in \mathbb{R}^n$. Also, this can be done globally, by calculating the masks using a globally calculated threshold, or locally to each layer, where a specific threshold is computed for each layer. When a set AN_t is calculated, it must be saved in the memory, because it will be used also for all the tasks $t' > t$, as well as during the inference phase. These sets are the only additional memory required by the method.

If the task is the first one, $t = 1$, the optimization problem to be solved is the following:

$$\min_{w, w_1} \mathbb{E}[\mathcal{L}(y, f(x; w \odot AN_1, w_1))]$$

where \odot is the element-wise product, and $w \odot AN_1$ means that, given the input sample x , and a layer l with the associated output z^l , the latter is masked using the binary mask present in AN_1 , accordingly to the layer's index, as: $\bar{z}^l = L_l(x) \odot m^{1,l}$. The masks

in AN_1 are used for zeroing the neurons which are not trained nor used during the training of the first task. In other words, it forces the minimization process to train only the neurons that were extracted using the scaling vectors associated to the current task, leaving the others unchanged.

If the task is not the first, $t > 1$, we want also to use neurons associated to past tasks, without changing them. To do this, for a task t and a layer l , we define the cumulative binary mask as:

$$c(t, l) = \prod_{i=1}^{t-1} m^{i,l}$$

where each mask $m^{i,l} \in AN_i$. Then, we define a set called Past Active Neurons for the task t as

$$PAN_t = \bigcup_{l=1}^L c(t, l).$$

This set contains, for each layer, all the cumulative binary masks $c(t, l)$. In other words, it contains all the neurons which were active during past tasks, and that must be preserved (not changed during the training), by masking the gradients associated to the weights which compose these neurons. At the same time, we want to use these neurons to improve the performance on the current task, by using the information they are capable of extracting.

Then, we need to extract the AN_t set, and this must be done by taking into consideration also the masks in PAN_t . It can be done in two different ways:

- **Soft extraction:** when calculating the threshold, also the neurons active in the PAN_t set are taken into consideration. Practically, if a past active neuron is chosen, nothing changes in the final result, since it belongs to PAN_t anyhow and it must not be changed in the minimization process. The dimension of the AN_t set is reduced, leaving more space for future tasks. This means that the procedure knows that some past neurons are more useful than others to solve the current task.
- **Hard extraction:** in this case, the threshold and the extracted neurons are calculated only over the neurons which are set to 0 in the PAN_t set.

Once we have these sets, we can minimize the optimization problem associated to the current task $t > 1$. Using the same notation as before, the optimization process for the task t is:

$$\min_{w, w_t} \mathbb{E}[\mathcal{L}(y, f(x; w \odot AN_t \odot PAN_t, w_t))]$$

To prevent the changing of the weights used by past tasks, we set to zero the gradients associated to these weights, and this is done at each step by multiplying these before the optimization step, as:

$$w' = w - \eta \left[\frac{\partial \mathcal{L}(\cdot, \cdot)}{\partial w} \odot (\sim PAN_t) \right]$$

where w' are the new weights, η is the learning rate, and \sim is the not operator, that negates all the masks in PAN_t ($\sim m^{t,l} = 1 - m^{t,l}$). The overall procedure is displayed in Fig. 2.

When we need to classify a sample x from a task t , we follow the same steps exposed before: the AN_t set is retrieved, if the task is not the first one PAN_t is created, then, we use these sets to mask out the outputs of the layers while performing the forward step, using the function $f(x; w \odot AN_t \odot PAN_t, w_t)$.

4. Experiments

In Section 4.1 we compare Structured Ensemble with different established ensemble methods, using multiple architectures and classification datasets. We evaluate different aspects of the ensemble models, such as accuracy and memory footprint 4.1.1, calibration 4.1.2, and the capacity to detect out-of-distribution samples 4.1.3 and 4.1.4. In the end, different ablation studies are shown in Section 4.1.5, in order to understand how the percentage of pruned neurons affects the final results.

In Section 4.2 we show the results on multiple CL benchmarks from the point of view of accuracy 4.2.1 and memory 4.2.2 required by each method. We conclude with a number of ablation experiments 4.2.3, in which we describe how the various hyper-parameters affect the final result on the CL benchmarks.

The code and all the files used to run the experiments are available online.¹

4.1. Classification

Datasets and architectures: we evaluate Structured Ensemble on different supervised classification datasets: CIFAR10 and CIFAR100 (Krizhevsky, Hinton, et al., 2009), SVHN (Netzer et al., 2011), and TinyImagenet (a subset of ImageNet (Deng et al., 2009) that contains 200 classes). We test our method on different architectures, to study how the extraction of the sub-networks behaves. The architectures are ResNet-20 (He, Zhang, Ren, & Sun, 2016) and VGG11 (Simonyan & Zisserman, 2014) for, respectively, CIFAR10-100 and SVHN, while for TinyImagenet we used ResNet-30 (He et al., 2016) and VGG16 (Simonyan & Zisserman, 2014). In this way we can test the extraction on sequential large networks such as VGG, and others with fewer parameters and with non-sequential architectures (due to the residual connections), like ResNet. Regarding the latter, applying the scaling vectors on each layer is not possible, due to the residual connection. For this reason we apply the scaling vectors on all the layers, except for the last one in each residual block. In each experiment we use 5 models to build each ensemble model.

Baselines: being our method an approach which aims to reduce the memory footprint of an ensemble model, we compare it with different established approaches that decrease the required memory or reduce the training time. For these reason we decided to compare it with **BatchEnsemble** (Wen et al., 2020), **Snapshot Ensemble** (Huang et al., 2017), **MC-Dropout** (Gal & Ghahramani, 2016), besides the **Single Model** approach (training a single neural network) and the standard ensemble of neural networks, which we call **Deep Ensemble**. These last two approaches are used only to provide a lower and upper baseline in terms of performance. We use the whole dataset for training each network, without performing any kind of bootstrapping.

Hyper-parameters: to set the hyper-parameters for each method we followed the respective papers. Regarding our method we test different percentages of pruning, and all obtained results are exposed in Section 4.1.5. The best trade-off is obtained with hard extraction done layer-wise with percentage is equal to 50%, and we train the scaling vectors for 10 epochs at the beginning of each task, setting $\lambda = 0.1$.

Training procedure: for each dataset and network combination we use the same training approach: we train for 200 epochs using SGD optimizer with momentum equal to 0.9, and starting learning rate equal to 0.01 for VGG and 0.1 for ResNet, that is annealed by multiplying it for 0.8 every 50 epochs. The only exception is Snapshot Ensemble, that we train for 50 epochs per cycle. Also, we extract 10% of the training split to build up the

¹ <https://github.com/jaryP/StructuredEnsemble>.

Table 1

Classification accuracy in percentage, on the classification benchmarks, using different architectures. The values are averaged over 3 experiments, and the standard deviation is also shown. Best results (among the ensembling approaches) for each dataset and network are highlighted in **bold**. Except for MC-Dropout (for which we do 5 sampling while testing) and Single model, each method is an ensemble of 5 neural networks. For TinyImagenet we used ResNet32 and VGG16; for the others dataset we used ResNet20 and VGG11.

Network	Method	SVHN	CIFAR10	CIFAR100	TinyImagenet
ResNet	Single model	94.08±0.72	88.67±0.47	59.71±1.59	44.82±0.79
	Deep ensemble	96.24±0.11	92.35±0.05	70.08±0.67	53.96±0.10
	Snapshot	95.69±0.91	90.38±0.08	61.84±0.11	48.75±0.53
	Batch ensemble	94.57±0.72	88.09±0.57	58.82±1.46	46.61±1.92
	MC-Dropout	94.52±0.01	87.58±0.79	58.33±2.02	43.33±2.10
	Structured Ensemble (ours)	96.18 ±0.04	90.95 ±0.05	67.55 ±0.66	49.72 ±0.21
VGG	Single model	94.26±0.03	88.67±0.47	52.69±2.24	44.20±2.01
	Deep ensemble	95.80±0.06	92.35±0.10	60.18±0.29	52.45±0.60
	Snapshot	95.35±0.13	89.38±0.08	59.85±0.86	47.38±0.19
	Batch ensemble	94.46±0.50	88.09±0.57	55.78±1.21	48.45±0.11
	MC-Dropout	94.16±0.02	87.58±0.79	53.87±1.25	44.02±0.96
	Structured Ensemble (ours)	95.74 ±0.05	89.84 ±0.08	63.13 ±1.02	52.58 ±0.42




Fig. 3. The image displays the overhead required by each method with respect to the base model, which requires 1. The results are associated to an ensemble of 5 ResNet20.

evaluation dataset, and we use the accuracy score calculated on it to save the best model during the training, as well as the early stopping criterion (with tolerance 5 for SVHN and 20 for the other datasets).

We normalize each dataset so that the images pixels are in the range [0, 1]. Also, for CIFAR10-100 and SVHN we use the standard augmentation scheme, in which the images are zero-padded with 4 pixels on each side and randomly cropped to produce 32×32 images, then horizontally mirrored with probability 0.5; we use the same scheme also for SVHN and TinyImagenet, by scaling the cropped dimensions of the latter according to the sizes of its images.

We repeat each experiment 3 times, using incremental seeds from 0 to 2.

4.1.1. Accuracy and memory

The accuracy results are summarized in Table 1, while Fig. 3 displays memory usage for each method in terms of the number of trainable parameters of the method over the number of trainable parameters required by a single neural network, called Overhead (e.g. a deep ensemble model composed by 5 networks has overheard equal to 5).

We can see, as expected, that all the models obtain a better, or similar, accuracy score than the single network, but worse than the naive ensemble approach. The methods that emulate an ensemble, Batch Ensemble and MC-Dropout, give the worst accuracy results, but require a small amount of memory (respectively close or equal to a single model). On the other hand, Snapshot and Structured Ensemble require more memory but the achieved

Table 2

The resulting ECE (lower is better) scores obtained for all the methods using ResNet architectures (results associated to the ones displayed in Table 1). The best results are highlighted in bold.

	CIFAR10	CIFAR100	TinyImagenet
Single model	8.21±0.17	19.60±3.97	3.19±0.79
Deep ensemble	1.61±0.04	7.69±0.32	1.56±0.72
Snapshot	3.01 ±0.21	12.93±0.43	5.60±0.12
Batch ensemble	8.14±1.24	20.52±1.54	2.73±0.14
MC-Dropout	7.11±0.12	13.63±1.83	22.54±0.87
Structured Ensemble (ours)	3.91±0.30	4.02 ±0.23	2.24 ±0.33

scores are close, even higher when using VGG on CIFAR100 or TinyImagenet, to the naive ensemble approach. In terms of memory, there is no difference between Snapshot Ensemble and the Deep Ensemble approach, while the memory used by Structured Ensemble depends on the pruning percentage used (as we will see).

4.1.2. Calibration

The calibration of a network is well defined and can be easily calculated. Given a sample x , the associated ground truth label y , the predicted class \hat{y} and its associated probability \hat{p} , we define the calibration of a method as:

$$p(\hat{y} = y \mid \hat{p} = p) = p \quad \forall p \in [0, 1].$$

This quantity cannot be computed with a finite set of samples, but, as proposed in Niculescu-Mizil and Caruana (2005), it can be approximated using a metric called Expected Calibration Score (ECE, lower is better), that measures how much the classification's probabilities over a dataset are miscalibrated, by combining two different quantities. Before defining how to calculate the ECE, we need to choose the number Z of equally sized bins in the range [0, 1], and grouping a set of predictions into Z interval bins B_z . Each B_z groups together the samples with a prediction confidence \hat{p} that falls into the range $(\frac{z-1}{Z}, \frac{z}{Z}]$. The ECE metric is calculated as:

$$ECE = \sum_{z=1}^Z \frac{|B_z|}{n} \left| \text{acc}(B_z) - \text{conf}(B_z) \right|$$

where n is the number of samples, $\text{acc}(B_z)$ and $\text{conf}(B_z)$ are, respectively, the averaged accuracy and the averaged confidence of the samples in B_z .

In Table 2 the resulting ECEs associated to ResNet20 trained on CIFAR10 are displayed. The scores obtained depends heavily on the dataset type. For CIFAR10 we have that the best method is Snapshot Ensemble, and Structured Ensemble is the second best;

Table 3

The diversity metrics obtained using all the methods with ResNet20 trained on CIFAR10 (results associated to the ones displayed in Table 1).

	CC diversity	WC diversity
Deep ensemble	5.73±0.42	32.89±1.56
Snapshot	5.60±0.12	31.67±0.77
Batch ensemble	2.73±0.14	19.21±1.92
MC-Dropout	3.62±0.45	22.54±0.87
Structured Ensemble (ours)	3.99±0.70	26.61±1.45

the other methods achieve a result that is similar to the Single Model, which is poorly calibrated. When the dataset becomes more difficult, Snapshot ensemble is one of the worst methods; we believe that this happens because the length of the learning cycles must be set accordingly to the dataset. On all the other dataset our proposal is the method that achieves the better ECE score. In the end, Structured Ensemble is the approach that gives the most reliable results.

4.1.3. Diversity analysis

The evaluation metrics used, such as accuracy and ECE score, do not express the key property of an ensemble model: the diversity of the models that compose it. In Fort et al. (2019) the authors proposed to measure the diversity of two networks based on how many samples they classify in the same way. We use a slightly modified version, that measures the similarity using the probabilities predicted by each model in the ensemble.

We define the function $g_i(x)$, which is associated to a network i , and, for a given sample x , it returns the unbounded prediction vector $v \in \mathbb{R}^c$, where c are the classes of the dataset taken into consideration. Give this formulation, we define the diversity calculated over a set of samples S , as $\frac{1}{S} \sum_i^S H\left(\frac{1}{N} \sum_j^N g_j(x_i)\right)$, where $H(\cdot)$ is the function that transforms the input vector into a distribution using the soft-max function, and calculates the entropy of it, normalized in $[0, 100]$. Using this formulation we have that, for a single sample, the higher the entropy is, the more dissimilar networks' prediction are.

To evaluate the diversity of networks that compose an ensemble model, we calculate it on Correctly Classified (CC Diversity) samples and Wrongly Classified (WC Diversity), both calculated on the test set. In Table 3 the results obtained using an ensemble of 5 ResNet20 trained on CIFAR10 are displayed.

We can see that, overall, each method gives low diversity for CC samples, and higher for WC ones. Also, Batch Ensemble and MC-Dropout have smaller difference between CC and WC samples, while Deep Ensemble and Snapshot ensemble display wider gap between CC and WC samples. Regarding Structured Ensemble, we see that the predictions are sufficiently diverse, and the scores are similar to the one obtained by Deep Ensemble.

4.1.4. Uncertainty

Here we study how the methods behave when it comes to detect Out-of-Distribution (OOD) samples. Our main idea is to study if a method can detect and discard less confident predictions, improving the accuracy (if compared to the same prediction without discarding samples).

To do this we rely on two different experiments. In the first one we follow the same scenario proposed in Pomponi et al. (2021): we attack the ensemble models using Fast Gradient Sign Method (Goodfellow, Shlens, & Szegedy, 2015) that, given an image x and its label y , corrupt x as $x + \epsilon \text{sign}(\nabla_x J(\theta, x, y))$, where J is the input-output Jacobian and θ is the set that contains the parameters of the network. The idea is to understand if the model is capable of detecting whether an image has been attacked too much, making correct classification impossible. In the second

experiment we evaluate the uncertainty on the recently proposed Corrupted CIFAR10 dataset (Hendrycks & Dietterich, 2019). The dataset consists of the same 10000 images present in the test spit of CIFAR10, but corrupted with 20 types of corruptions, each one with 5 different severity levels.

To understand if an image is corrupt and must be discarded we use the discarding process proposed in Pomponi et al. (2021). After a preliminary evaluation of the method, we decided to use the following formulation:

$$T_\gamma(\mathcal{H}) = Q_{\mathcal{H}}(75)$$

where function $Q_{\mathcal{H}}(75)$ returns the percentile associated to the 75%-th value of the set \mathcal{H} . In our experiments \mathcal{H} is the entropy of each correctly classified image in the development set, calculated after the training process.

Given a threshold and a corrupted set of samples that we want to filter by discarding highly corrupted images, we evaluate the performances by measuring different quantities:

- Accuracy (A): the accuracy calculated over the corrupted set.
- Filtered Accuracy (FA): the accuracy calculated only over the images that have not been discarded.
- Discarded samples (D): the percentage of images that have been discarded.

A good method should be able to increment the accuracy by discarding images on which the confidence is low (maximize the difference between FA and A), while keeping the highest number of images on which the accuracy is calculated (minimize D).

In Table 4 the results obtained using ResNet20 trained on CIFAR10 are exposed. To study how the performances are affected by the corruptions, we average the scores obtained on moderate and high corrupted images, by setting ϵ equal to 0.02 and 0.5 for FGSM, and using severity equal to 3 and 5 for C-CIFAR10.

We see that Batch Ensemble and Snapshot are the worst methods. In fact, the latter is not capable of discarding enough corrupted images (due to the similarity between the networks), while the first discards more samples than the others methods, but is not capable of reaching good scores if compared to the other methods; this means that the discarded samples also contain images that would have been classified correctly. These results are expected, because these methods are not capable of building an ensemble in which the models are different enough. Regarding MC-Dropout, despite being proposed mainly for uncertainty estimation, it has lower performance than Structured Ensemble, especially when the severity value is high, where it reaches a score that is 4% higher than the second best one, while discarding a similar number of images.

4.1.5. Ablation studies

In this section we study how the pruning percentage affects the results. Table 5 contains all the results associated to the different pruning percentages studied.

Accuracy and overhead: from the table we can see that the accuracy decreases with the increase of pruned neurons, as expected, but it is always higher than the score obtained using a single network. Regarding the overhead, we can see that the additional memory required by the method is always lower than deep ensemble (which requires 5 in this case), and when the pruning percentage grows, Structured Ensemble produces an ensemble model that requires less memory than the base model, but achieves higher scores.

Calibration: the calibration of our proposal depends on the pruning percentage: we see that the calibration is poor when the pruning percentage is low, but, if we can sacrifice some accuracy points (around 2%), we can build up an ensemble which is better

Table 4

The results of the uncertainty used to discard the corrupted images. The results are associated to an ensemble model composed of 5 ResNet20 trained on CIFAR10, using the trained networks obtained in Table 1. The results are computed over the test split, which is composed by 10000 images. We only show the results associated to a medium and high corruption of the images (the results associated to Corrupted CIFAR10 are averaged over all the types of corruption). The format of the results is the following $A \xrightarrow{D} FA$, where the definition of A, D, and FA can be found in Section 4.1.4. This formulation means that the accuracy changes from A to FA by discarding D percentages of images with high associated uncertainty. For each corruption level we highlighted in bold the best Filtered Accuracy FA.

Method	FGSM		C-CIFAR10	
	ϵ		Severity	
	0.02	0.5	3	5
Snapshot	49.95 $\xrightarrow{45.62}$ 65.77	21.18 $\xrightarrow{48.07}$ 25.91	68.25 $\xrightarrow{52.53}$ 86.12	50.12 $\xrightarrow{57.34}$ 71.02
Batch ensemble	42.00 $\xrightarrow{69.08}$ 65.39	21.17 $\xrightarrow{80.13}$ 32.10	67.49 $\xrightarrow{54.45}$ 87.37	49.49 $\xrightarrow{69.97}$ 72.12
MC-Dropout	48.39 $\xrightarrow{71.70}$ 76.89	22.06 $\xrightarrow{82.51}$ 38.42	70.32 $\xrightarrow{52.83}$ 91.99	52.76 $\xrightarrow{67.36}$ 78.17
Structured Ensemble (ours)	48.06 $\xrightarrow{72.27}$ 78.57	21.08 $\xrightarrow{86.56}$ 41.88	70.01 $\xrightarrow{57.21}$ 92.10	51.27 $\xrightarrow{70.78}$ 82.12

Table 5

The resulting metrics obtained when changing the pruning percentage, averaged over 3 experiments (the standard deviation is also shown). The definitions of these metrics are in Section 4.1. The displayed results are associated to an ensemble model composed of 5 ResNet20 trained on CIFAR10.

	Accuracy	ECE	Overhead	CC diversity	WC diversity
30%	91.51±0.10	4.54±0.26	3.39±0.01	2.85±0.41	22.62±1.07
50%	90.66±0.12	3.91±0.30	2.39±0.01	3.99±0.70	26.61±1.45
60%	90.59±0.46	3.61±0.23	2.38±0.02	3.88±0.36	25.92±1.76
70%	89.43±0.32	2.78±0.25	1.42±0.04	5.87±0.26	31.12±0.26
80%	88.40±0.51	1.55±0.24	0.92±0.02	8.21±0.27	35.42±0.50

Table 6

The resulting metrics obtained when changing the initialization of the masks, using different distributions from which the values are sampled. The definitions of these metrics are in Section 4.1. The displayed results are associated to an ensemble model composed of 5 ResNet20 trained on CIFAR10, build using Structured Ensemble with pruning percentage set to 50%.

	Accuracy	ECE	Overhead	CC diversity	WC diversity
$\mathcal{N}(0, 1)$	91.05	3.80	2.39	3.34	24.22
$\mathcal{N}(0, 2)$	90.93	4.11	2.39	3.34	24.22
$\mathcal{N}(0, 5)$	90.87	3.92	2.48	4.16	26.02
$\mathcal{U}(0, 1)$	90.86	4.01	2.41	5.47	25.03
$\mathcal{U}(-1, 1)$	90.74	3.94	2.30	4.74	29.11

calibrated than all the other methods. These results are expected, because the accuracy score and the ECE are deeply connected, and usually, a network loses the calibration of its predictions with the improving of the accuracy. In fact, ensembles composed by smaller networks give slightly worst results (around 1–3% of accuracy is lost) but better ECE.

Diversity: from the table we see that the CC diversity is always low regardless of the pruning percentage, but WC diversity depends on the pruned neurons, and it grows with the growing of the pruning percentage. This phenomenon can be explained thinking that smaller networks tend to be very different from each other, while bigger ones can converge to similar minima, and may give similar predictions. In fact, with the increase of the pruning percentage also the CC diversity increases, and with it the difference between CC and WC diversities.

Initialization of the masks: in Table 6 the results obtained when varying the initialization of the masks are shown. We can see that the obtained results do not depend heavily on the initialization of the masks, because these are used mainly to estimate the gradient value associated to each output neuron. However, when the variance becomes too large, we observe a slight decrease of the performances. In the end, we decided to initialize each mask using the standard distribution $\mathcal{N}(0, 1)$.

4.2. Continual learning

Continual learning setup: any dataset can be used to build a CL dataset. To do this, we group different labels from the dataset into disjoint subsets, and, for each set of labels, we extract the samples associated to its components from the original dataset. In this way, we can split a dataset having c classes into M tasks, each one composed by $\frac{c}{M}$ classes. We use this established approach to build 3 CL benchmarks: (1) Split-MNIST, which contains 5 tasks, each one with two classes, (2) Split-CIFAR10 (Rebuffi, Kolesnikov, Sperl, & Lampert, 2017), with the same configuration of Split-MNIST, and (3) Split-CIFAR100, that contains 10 tasks, each one with 10 classes. For MNIST we used LeNet-5 (LeCun et al., 1989), while for the others datasets we used VGG11 (Simonyan & Zisserman, 2014), with the exception that we halved the size of each kernel. We trained all networks using Adam optimizer (Kingma & Ba, 2014), with learning rate equal to 0.001. To have a better statistic of the results we repeat each experiment 3 times, as in the ensemble experiments.

Baselines: being our method a structural one, we compare it with similar approaches proposed in the recent years: **BatchEnsemble** (Wen et al., 2020) (also used in the ensemble experiments), **Supermask in Superposition** (Wortsman et al., 2020), and **Pruning** (Golkar et al., 2019). Also, we selected two other baselines, which are not CL methods: the first one is **Single Model** approach, in which we train a model for each task, and the second is **Naive**, where a single network is trained sequentially on the tasks without any method to alleviate CF. These last two can be seen as, respectively, upper and lower bounds for the scores achievable by the other methods; for this reason we do not compare the CL methods directly to them.

Hyper-parameters: to find the best hyper-parameters for each method we followed the respective paper. Regarding our method we test different combination of parameters, which are the same exposed before. We ended up with hard extraction done layer wise with percentage is equal to 50%, and we train the scaling vectors for 10 epochs at the beginning of each task.

Metrics: to evaluate the efficiency and to compare the CL methods usually two metrics are used. One shows the effective accuracy obtained on all the tasks, while the second indicates the amount of accuracy over past tasks that is lost while training new ones. In this case, the selected CL methods completely remove the CF phenomenon and for this reason we use only one metric from (Díaz-Rodríguez, Lomonaco, Filliat, & Maltoni, 2018), simply called Accuracy. It is calculated on a matrix $R \in \mathbb{R}^{M \times M}$, where M is the number of tasks, and each entry R_{ij} is the test accuracy on task j when the training on task i is completed. Given this matrix of scores, the accuracy is calculated as:

$$\text{Accuracy} = \frac{\sum_{i>j}^M R_{ij}}{\frac{1}{2}M(M+1)}.$$

Table 7

Final Accuracy averaged over all the tasks, evaluated when the last one is over, is shown. The values are averaged over 3 experiments, and the standard deviation is also shown. Best results for each dataset are highlighted in **bold**. Being the all methods structural, the CF phenomenon is not present (it is present only using the naive approach, but it is not shown).

Method	S-MNIST	S-CIFAR10	S-CIFAR100
Naive	94.32±0.04	69.99±2.45	29.39±0.40
Single model	99.85±0.04	92.62±0.82	70.44±0.15
SupSup	99.45±0.02	91.06±0.11	65.59±0.34
Pruning	97.91±0.76	78.23±1.56	49.27±0.49
Batch ensemble	99.34±0.07	85.40±0.07	56.12±0.30
Structured Ensemble (ours)	99.66±0.26	92.27±0.32	66.24±0.39

Table 8

The overhead, in terms of floats F (e.g. trainable parameters) and binaries B (e.g. binary values) in addition to the parameters of the neural network. These values are an estimation, and are related to the additional values stored within the model once the training phase is over. For comparison, the number of floats parameters of a base network is indicated with the constant b , and the number of tasks is M .

Method	F	B
Single model	$(M - 1) \cdot b$	0
SupSup	0	$M \cdot b$
Pruning	0	$M \cdot b$
Batch ensemble	$< 0.1 \cdot b \cdot M$	0
Structured Ensemble (ours)	0	$< 0.001 \cdot b \cdot M$

Table 9

The number of bytes used by each method for two networks. The Naive approach is used to show the number of base memory required by each model, which is b . The additional memory required by each method depends on the number of tasks M . The results are correlated with the ones exposed in Table 8. Here, we set $F = 64$, because we assumed that a float requires 64 bites to be saved in the memory, and $B = 1$, because each binary value requires 1 bit.

Method	LeNet-5	half VGG11
Naive (b)	50550 · 64	2304864 · 64
Single model	50550 · 64 · M	2304864 · 64 · M
SupSup	50550 · M	2304864 · M
Pruning	50550 · M	2304864 · M
Batch ensemble	164 · 64 · M	2499 · 64 · M
Structured Ensemble (ours)	142 · M	1376 · M

In addition to this metric, we compared the methods under the point of view of the memory required, by approximating the number of additional float or binary values used.

4.2.1. Classification results

In Table 7 the results are shown. We can see that our method outperforms the others, since it is capable of achieving higher performances (while also requiring small additional memory, as we will see). This is accentuated when the complexity of the problem increases, starting from CIFAR10, in which our method gives performances which are similar to the upper bound (Single Model baseline). We can conclude that our intuition about the usability of past information improves the performances, compared to the other methods that use a more neat separation of the sub-networks for solving the tasks.

4.2.2. Memory usage

Memory is a crucial aspect of any CL approach, because the number of tasks can grow over time, and so it is important to keep the required memory limited. Table 8 approximates the memory required by each method, by counting the number of additional values stored in the memory while taking into account the type of the number to save. We partition memory into two components: the floats (F), and the binaries (B). The first one counts the additional values that must be stored as float values (e.g. the vectors in Batch Ensemble), while the second one counts

Table 10

The results obtained on S-CIFAR10 scenario by combining pruning techniques, which are: Hard (H) and Soft (S) extraction. Each one can be done locally in each layer (L) or globally (G). The results are averaged over two experiments, and best one is highlighted in bold.

L-S	92.14±0.13
L-H	92.50±0.11
G-S	91.90±0.31
G-H	92.26±0.22

only the values that can be saved as 0 or 1. The split is necessary because a binary value requires less storage memory than a float (usually $B = 64F$). Between the methods, the worst are the one that apply binary masks over the whole networks (SupSup and Pruning), followed by BatchEnsemble and Structured Ensemble, with the latter being capable of further reducing memory usage, making it negligible.

In Table 9 we expose the approximated count of additional bits used by each method. We can see that our proposal is the best method, as said before. In fact, taking LeNet-5 as reference, we have that our method can store 355 ($\frac{50550}{142}$) masks, used to solve the corresponding number of tasks, using the same space that SupSup uses to store the mask for a single task. The second best method in terms of memory, BatchEnsemble, can store only 4 tasks using the same space (due to the float values used by the method).

In the end, the best methods are Batch Ensemble and our proposal, Structured Ensemble. Both of the approaches require small memory overhead (with the latter being better under this aspect). Taking in consideration both the accuracy results exposed before (Table 1) and the memory Tables 8 and 9, we conclude that our method is the best in terms of both memory requirements and accuracy obtained.

4.2.3. Hyper-parameters of structured ensemble

In this section we study how the hyper-parameters of our method affect the results obtained on S-CIFAR10 dataset.

Local vs Global pruning: in Table 10 the results of the various combinations are shown. The results show clearly that the local approach is always better than the global one. To understand why it happens we inspected the scaling vectors produced by the method, and discovered that when using the global approach, the method tends to assign all the neurons in the first layers to the first task, because these are the ones which extract the features directly from the images. This prevents future tasks to adjust how these features are extracted from the images, and so they are forced to work with inappropriate high level features. On the contrary, using the local approach, a new task can always change the inactive neurons of the first layers, allowing custom high level features extraction.

Hard vs Soft extraction: Table 10 also shown the difference between soft and hard extraction. In the local approach, we achieve better results using hard extraction approach, because using this schema the method is always capable of allocating new learnable neurons for the current task, while in the soft pruning case the number of active neurons may vary, and this could lead to a lack of neurons that can be used to solve properly the current task. This phenomenon is also visible when it comes to the global pruning, but the scores are not good enough in either case, due to the limitations exposed before.

Pruning percentage: Fig. 4 shows how the accuracy score varies with the pruning percentage (the results are obtained using local-hard extraction). The pruning percentage must be set in such a way that the first task has all the computational capacity to be solved correctly. At the same time, this must be achieved




Fig. 4. How pruning percentage affects the final score obtained on S-CIFAR10, with L-H setting.

using the smallest number of neurons, in order to leave some of them free for future tasks.

We can see that the worst results are obtained when the pruning percentage is too low, around 10%, or too high, around 90%, because, respectively, the space allocated for future tasks is small or the first one has a small amount of network to use (the first case is slightly preferable, because at least the first task can be solved efficiently, and this can also improve the accuracy on future tasks). The best values are achieved with a percentage in [50, 55]%, which is the final value used in all the experiments exposed before.

5. Conclusion

In this paper we proposed a novel approach to build ensembles of deep models requiring less memory. Our approach, called Structured Ensemble, is capable of extracting multiple networks from a single one, and use these to build a smaller ensemble model, which gives comparable results. Each network is different from the others, and contains the minimal set of structural connections from the original network that can be trained to achieve good accuracy while reducing the required memory.

We evaluated our proposal on a wide set of different experiments, including an evaluation of their robustness and diversity, showing it performs on-par or better than state-of-the-art approaches in the ensemble literature. We also evaluated our method in a continual learning scenario, in which the model should be capable of solving a stream of tasks without losing the ability to solve past learned ones. Here too we compared our method to similar approaches, on a wide set of benchmarks, and we achieved the best scores while keeping the memory requirement negligible.

Future work will explore the time aspect of Structured Ensemble, by focusing on techniques that can speed up the training process, while keeping the ability to extract multiple models from a single one, keeping the memory requirements contained.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

Allen-Zhu, Z., & Li, Y. (2020). Towards understanding ensemble, knowledge distillation and self-distillation in deep learning. *arXiv:2012.09816*.
 Breiman, L. (1996). Bagging predictors. *Machine Learning*, 24(2), 123–140.
 Bucilua, C., Caruana, R., & Niculescu-Mizil, A. (2006). Model compression. In *Proceedings of the 12th ACM SIGKDD international conference on knowledge discovery and data mining* (pp. 535–541).
 Deng, J., Dong, W., Socher, R., Jia Li, L., Li, K., & Fei-fei, L. (2009). Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition* (pp. 248–255).

Díaz-Rodríguez, N., Lomonaco, V., Filliat, D., & Maltoni, D. (2018). Don't forget, there is more than forgetting: new metrics for continual learning. In *Workshop on continual learning, NeurIPS 2018 Neural Information Processing Systems* (pp. 1–7).
 Fort, S., Hu, H., & Lakshminarayanan, B. (2019). Deep ensembles: A loss landscape perspective. *ArXiv preprint arXiv:2012.09816*.
 Frankle, J., & Carbin, M. (2018). The lottery ticket hypothesis: Finding sparse, trainable neural networks. In *International conference on learning representations* (pp. 1–42).
 French, R. M. (1999). Catastrophic forgetting in connectionist networks. *Trends in Cognitive Sciences*, 3(4), 128–135.
 Gal, Y., & Ghahramani, Z. (2016). Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. In *International conference on machine learning* (pp. 1050–1059). PMLR.
 Garipov, T., Izmailov, P., Podoprikin, D., Vetrov, D., & Wilson, A. G. (2018). Loss surfaces, mode connectivity, and fast ensembling of DNNs. In *Proceedings of the 32nd international conference on neural information processing systems* (pp. 8803–8812).
 Golkar, S., Kagan, M., & Cho, K. (2019). Continual learning via neural pruning. *ArXiv preprint arXiv:1903.04476*.
 Goodfellow, I., Shlens, J., & Szegedy, C. (2015). Explaining and harnessing adversarial examples. In *International conference on learning representations* (pp. 1–11).
 Gretton, A., Borgwardt, K. M., Rasch, M. J., Schölkopf, B., & Smola, A. (2012). A kernel two-sample test. *Journal of Machine Learning Research*, 13, 723–773.
 Guo, C., Pleiss, G., Sun, Y., & Weinberger, K. Q. (2017). On calibration of modern neural networks. In *International conference on machine learning* (pp. 1321–1330). PMLR.
 Gustafsson, F. K., Danelljan, M., & Schon, T. B. (2020). Evaluating scalable Bayesian deep learning methods for robust computer vision. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops* (pp. 318–319).
 He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 770–778).
 He, Y., Zhang, X., & Sun, J. (2017). Channel pruning for accelerating very deep neural networks. In *Proceedings of the IEEE international conference on computer vision* (pp. 1389–1397).
 Hendrycks, D., & Dietterich, T. (2019). Benchmarking neural network robustness to common corruptions and perturbations. *ArXiv preprint arXiv:1412.6572*.
 Hinton, G., Vinyals, O., & Dean, J. (2015). Distilling the knowledge in a neural network. In *NIPS deep learning and representation learning workshop* (pp. 1–9).
 Huang, G., Li, Y., Pleiss, G., Liu, Z., Hopcroft, J. E., & Weinberger, K. Q. (2017). Snapshot ensembles: Train 1, get m for free. *ArXiv preprint arXiv:1704.00109*.
 Kawaguchi, K. (2016). Deep learning without poor local minima. In *Advances in neural information processing systems* (Vol. 29) (pp. 1–9). Curran Associates, Inc..
 Kingma, D. P., & Ba, J. (2014). Adam: A method for stochastic optimization. *ArXiv preprint arXiv:1412.6980*.
 Krizhevsky, A., Hinton, G., et al. (2009). *Learning multiple layers of features from tiny images*. Citeseer.
 Krogh, A., & Vedelsby, J. (1995). Validation, and active learning. *Advances in Neural Information Processing Systems*, 7, 231.
 Kwon, Y., Won, J.-H., Kim, B. J., & Paik, M. C. (2020). Uncertainty quantification using Bayesian neural networks in classification: Application to biomedical image segmentation. *Computational Statistics & Data Analysis*, 142, Article 106816.
 Lakshminarayanan, B., Pritzel, A., & Blundell, C. (2017). Simple and scalable predictive uncertainty estimation using deep ensembles. In *Proceedings of the 31st international conference on neural information processing systems* (pp. 6405–6416).
 LeCun, Y., Boser, B., Denker, J. S., Henderson, D., Howard, R. E., Hubbard, W., et al. (1989). Backpropagation applied to handwritten zip code recognition. *Neural Computation*, 1(4), 541–551.
 Lee, N., Ajanthan, T., & Torr, P. (2018). Snip: Single-shot network pruning based on connection sensitivity. In *International conference on learning representations* (pp. 1–15).
 Lee, S., Purushwalkam, S., Cogswell, M., Crandall, D., & Batra, D. (2015). Why m heads are better than one: Training a diverse ensemble of deep networks. *ArXiv preprint arXiv:1511.06314*.
 Li, H., Kadav, A., Durdanovic, I., Samet, H., & Graf, H. P. (2016). Pruning filters for efficient convnets. *ArXiv preprint arXiv:1608.08710*.
 Lopez-Paz, D., & Ranzato, M. (2017). Gradient episodic memory for continual learning. In *Advances in neural information processing systems* (pp. 6467–6476).
 McCloskey, M., & Cohen, N. J. (1989). Catastrophic interference in connectionist networks: The sequential learning problem. In *Psychology of learning and motivation* (Vol. 24) (pp. 109–165). Elsevier.
 Netzer, Y., Wang, T., Coates, A., Bissacco, A., Wu, B., & Ng, A. Y. (2011). Reading digits in natural images with unsupervised feature learning.

- Niculescu-Mizil, A., & Caruana, R. (2005). Predicting good probabilities with supervised learning. In *Proceedings of the 22nd international conference on machine learning* (pp. 625–632). Association for Computing Machinery.
- Ovadia, Y., Fertig, E., Ren, J., Nado, Z., Sculley, D., Nowozin, S., et al. (2019). Can you trust your model's uncertainty? Evaluating predictive uncertainty under dataset shift. In *Advances in Neural Information Processing Systems (Vol. 32)* (pp. 1–25).
- Parisi, G. I., Kemker, R., Part, J. L., Kanan, C., & Wermter, S. (2019). Continual lifelong learning with neural networks: A review. *Neural Networks*, 113, 54–71.
- Perrone, M. P., & Cooper, L. N. (1992). *When networks disagree: Ensemble methods for hybrid neural networks*. Brown University, institute for brain and neural system.
- Pomponi, J., Scardapane, S., & Uncini, A. (2021). Bayesian neural networks with maximum mean discrepancy regularization. *Neurocomputing*.
- Ramamujan, V., Wortsman, M., Kembhavi, A., Farhadi, A., & Rastegari, M. (2020). What's hidden in a randomly weighted neural network? In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 11893–11902).
- Rebuffi, S.-A., Kolesnikov, A., Sperl, G., & Lampert, C. H. (2017). iCARL: Incremental classifier and representation learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 2001–2010).
- Rusu, A. A., Rabinowitz, N. C., Desjardins, G., Soyer, H., Kirkpatrick, J., Kavukcuoglu, K., et al. (2016). Progressive neural networks. ArXiv preprint [arXiv:1606.04671](https://arxiv.org/abs/1606.04671).
- Simonyan, K., & Zisserman, A. (2014). Very deep convolutional networks for large-scale image recognition. ArXiv preprint [arXiv:1409.1556](https://arxiv.org/abs/1409.1556).
- Thrun, S. (1998). Lifelong learning algorithms. In *Learning to learn* (pp. 181–209). Springer.
- Van de Ven, G. M., & Tolias, A. S. (2019). Three scenarios for continual learning. ArXiv preprint [arXiv:1904.07734](https://arxiv.org/abs/1904.07734).
- Wen, Y., Tran, D., & Ba, J. (2020). BatchEnsemble: an alternative approach to efficient ensemble and lifelong learning. In *International conference on learning representations* (pp. 1–20).
- Wortsman, M., Ramamujan, V., Liu, R., Kembhavi, A., Rastegari, M., Yosinski, J., et al. (2020). Supermasks in superposition. *Advances in Neural Information Processing Systems*, 33.