ORIGINAL ARTICLE

Expert Systems | WILEY

Efficient adaptive ensembling for image classification

Bruno Antonio [ID] | Davide Moroni [ID] | Massimo Martinelli [ID]

Institute of Information Science and Technologies, Italian National Research Council, Pisa, Italy

Correspondence
Massimo Martinelli, Institute of Information Science and Technologies, Italian National Research Council, Via Giuseppe Moruzzi 1, Pisa 56124, Italy.
Email: massimo.martinelli@isti.cnr.it

Abstract

In recent times, with the exception of sporadic cases, the trend in computer vision is to achieve minor improvements compared to considerable increases in complexity. To reverse this trend, we propose a novel method to boost image classification performances without increasing complexity. To this end, we revisited ensembling, a powerful approach, often not used properly due to its more complex nature and the training time, so as to make it feasible through a specific design choice. First, we trained two EfficientNet-b0 end-to-end models (known to be the architecture with the best overall accuracy/complexity trade-off for image classification) on disjoint subsets of data (i.e., bagging). Then, we made an efficient adaptive ensemble by performing fine-tuning of a trainable combination layer. In this way, we were able to outperform the state-of-the-art by an average of 0.5% on the accuracy, with restrained complexity both in terms of the number of parameters (by 5–60 times), and the FLoating point Operations Per Second FLOPS by 10–100 times on several major benchmark datasets.

KEYWORDS
convolutional neural networks, deep learning, EfficientNet, ensemble, image classification

1 | INTRODUCTION

Computer vision is one of the fields that most benefit from deep learning, continuously improving the state-of-the-art (SOTA) using convolutional neural networks (CNNs) and visual transformers. In nearly all computer vision scenarios, complexity grows exponentially, even for minimal improvements, both in terms of the number of parameters and in FLoating point Operations Per Second (FLOPS). Table 1 briefly shows the evolution of the SOTA on the ImageNet classification task. It can be observed that the trend of improvements achieved only through high complexity growth was temporarily slowed down by the introduction of EfficientNet architecture (and in particular with EfficientNet-b0 attaining the best accuracy/complexity trade-off; Tan & Le, 2019). This also applies to other image classification datasets (e.g., CIFAR) and to computer vision tasks based on CNNs (e.g., object detection and segmentation).

Among the various machine learning approaches, ensembling is a technique that combines several models, called weak learners, in order to produce a model with better performance than any of the weak learners alone (Opitz & Maclin, 1999). Usually, the combination is accomplished by aggregating the output of the weak learners, generally this is made by voting (resp. averaging) for classification (resp. regression). Other aspects, such as ensemble size (i.e., number of weak learners) and ensemble techniques (e.g., bagging, boosting, stacking), are crucial for obtaining a satisfactory result. Since it requires the training of several models, ensembles makes the overall validation much more expensive, and model complexity grows at least linearly compared to the ensemble size. Moreover, ensembling is a time-consuming process, and this is the main reason

TABLE 1 Evolution of the state-of-the-art on the ImageNet classification task: as can be seen, complexity in models having accuracy >80% (both in the number of parameters and FLOPs) grows exponentially despite the slightest improvement. The same trend can be noticed in other computer vision tasks. N.B. only some architectures providing relevant improvements are shown.

Model	Year	Accuracy	Parameters	FLOPs
AlexNet (Krizhevsky et al., 2012)	2012	63.3%	$\approx 60\,M$	$\approx 0.7\,G$
InceptionV3 (Szegedy et al., 2016)	2015	78.8%	$\approx 24\,M$	$\approx 6\,G$
ResNeXt-101 64×4 (Xie et al., 2017)	2016	80.9%	$\approx 84\,M$	$\approx 16\,G$
EfficientNet-b0 (Tan & Le, 2019)	2019	77.1%	$\approx 5.3\,M$	$\approx 0.4\,G$
EfficientNet-b7 (Tan & Le, 2019)	2019	84.3%	$\approx 67\,M$	$\approx 37\,G$
Swin-L (Liu et al., 2021)	2021	87.3%	$\approx 197\,M$	$\approx 103\,G$
NFNet-F4+ (Brock et al., 2021)	2021	89.2%	$\approx 527\,M$	$\approx 215\,G$
ViT-G/14 (Zhai et al., 2021)	2021	90.45%	$\approx 1843\,M$	$\approx 965\,G$
CoAtNet-7 (Dai et al., 2021)	2021	90.88%	$\approx 2440\,M$	$\approx 2586\,G$

preventing a more extended use in practice, especially in computer vision. On the contrary, this work shows that our technique exploits this powerful tool with limited resources (e.g., compared to the model complexity, validation time and training time).

This work shows how applying a well-defined ensembling strategy, using an efficient basic model as the core, can improve the state-of-the-art in computer vision tasks, preserving a competitive performance/complexity trade-off. In Section 3 we describe our design strategy in detail (e.g., model, ensembling strategy, validation), focusing on the introduction of the main novel aspects. Experimental results and data description are shown in Section 4, while an exhaustive discussion is provided in the last section.

2 | RELATED WORK

In recent years, the demand for intelligent systems based on image processing has also grown on the push of emerging business markets. In this context, the capacity to deal with large-scale collections of images has not only to face significant technological challenges but must be shown to be cost-effective and, ultimately, sustainable. Indeed, the carbon impact of artificial intelligence (AI) is a concern that has been well recognized (Dhar, 2020), favouring the adoption of green AI paradigms (Schwartz et al., 2020). In particular, in order to reduce the carbon footprint of AI and make it cost-effective in new markets, it is possible to follow several pathways, including decentralized approaches based on federated learning (therefore not requiring energy-consuming data transfer) (Bonawitz et al., 2019) or devising ad hoc low-consumption hardware specific for modern deep learning algorithms (Sze et al., 2017). Other methods deal with the AI model itself, proposing its simplification or optimization; well-known techniques, mainly suited for inference, include parameter quantization and pruning, compressed convolutional filters and matrix factorization, network architecture search, and knowledge distillation (Goel et al., 2020). In this paper, instead, we propose a method for achieving greener models both in training and inference by resorting to ensembling.

Ensembling consists in a machine learning approach in which a set of weak learners (or basic models) is turned into a strong learner (or ensemble model) (Opitz & Maclin, 1999; Sagi & Rokach, 2018). The set of weak learners might consist of homogenous models (i.e., they are all from the same family or architecture) or might be heterogeneous, that is, the basic models belong to different machine learning paradigms. The basic example is to put together multiple models trained for solving the same classification or regression task and then combine them in some fashion, for example, by performing majority voting in the case of classification or averaging in the case of regression. The scope of performing ensembling is generally related to the desire to reduce the bias or variance that affects a machine learning task (Dong et al., 2020). As it is well known, a low-complexity model might have a significant error in attaining adequate performance on a dataset, even during training. This is commonly due to the low representation capabilities of simple models that can only capture some of the complex patterns in the training datasets. Such error during training is referred to as the bias of the model. By converse, very complex models have many degrees of freedom to completely stick to the training dataset and convey excellent performance during training. However, they apprehend not only the relevant features of the problem but also learn unimportant features of the training dataset. This results in relatively inadequate performance during test and validation: the model needs to be more balanced to the training dataset and reach good general results, having scarce generalization capabilities. Such an issue is usually indicated as a high variance of the model. The three primary techniques for conducting ensembling are bagging, boosting, and stacking. In general, bagging decreases the variance among the weak classifiers, while boosting-based ensembles reduce bias and variance. Stacking is generally employed as a bias-reducing procedure. In more detail, the bagging technique involves partitioning the training datasets into distinct subsets based on specific criteria, such as equalizing class distributions within each subset. Subsequently, a weak classifier is trained using each subset of the training set. Ideally, these classifiers possess low bias on the training set but may exhibit high variance. The outputs of these individual classifiers are then combined through weighted voting or a weighted average using a specially designed layer. This fusion of weak classifiers

forms the strong classifier, which tends to have reduced variance. It is worth noting that the weak classifiers can be trained independently and in parallel. In boosting instead, weak classifiers are very simple and low complexity but are trained cleverly, for example, using cascading. Ultimately, stacking commonly involves the consideration of diverse weak learners with varying characteristics. The training process takes place concurrently, and a final amalgamation is achieved by training a meta-model that generates predictions based on the collective inputs from the various weak models. In general, all of these approaches have been used in conjunction with deep learning models. The review (Ganaie et al., 2021) presents some recent literature on the subject systematically. In this paper, we propose using bagging in an original way that allows us to obtain superior results with respect to the state of the art while decreasing the computational burden.

## 3 | EFFICIENT ADAPTIVE ENSEMBLING

3.1 | Efficiency

At the foundations of the efficiency of the proposed method lies the basic core model adopted in this work: EfficientNet (Tan & Le, 2019). As the name suggests, EfficientNet improves the classification quality with lower complexity compared to models having similar classification performances. This is possible since EfficientNet performs optimized network scaling, given a predefined complexity. As shown in Figure 1, in the CNN literature, there are three main types of scaling: depth scaling, width scaling and input scaling. Depth scaling consists in increasing the number of layers in the CNN; it is the most popular scaling method in the literature and allows detecting features at multiple levels of abstraction. Width scaling consists in increasing the number of convolutional kernels and parameters or channels, giving the model the capability to represent different features at the same level. Input scaling is represented by the increase in size/resolution of the input images, which allows for capturing additional details.

Each of these scaling's can be manually set or via a grid search. However, they increase the model complexity, usually exponentially, with tons of new parameters to tune and, after a certain level, scaling appears not to improve performances. The scaling method introduced in (Tan & Le, 2019) is named compound scaling. It suggests that the strategic execution of all scaling together provides better results because it is argued that they are dependent. Intuitively, they introduce the compound coefficient ϕ representing the total amount of resources available to the model and find the optimal scaling combination given such a constraint, following the rules in Equation 1. In this way, the total complexity of the network is approximately proportional to 2^{ϕ} (see the original paper for more details).

$$\text{depth: } d = \alpha^{\phi} \quad \text{width: } w = \beta^{\phi} \quad \text{resolution: } r = \gamma^{\phi}$$
$$\text{such that } \quad \alpha \cdot \beta^2 \cdot \gamma^2 \approx 2 \quad \text{and} \quad \alpha \geq 1, \beta \geq 1, \gamma \geq 1 \tag{1}$$

3.2 | Adaptivity

The adaptivity is given by the fact that the proposed ensembling is data-driven and not fixed as usual. The typical way of combining weak learners is to perform voting/averaging as shown in Figure 2 (predicting the output from all weak learners and then picking the most frequent output/
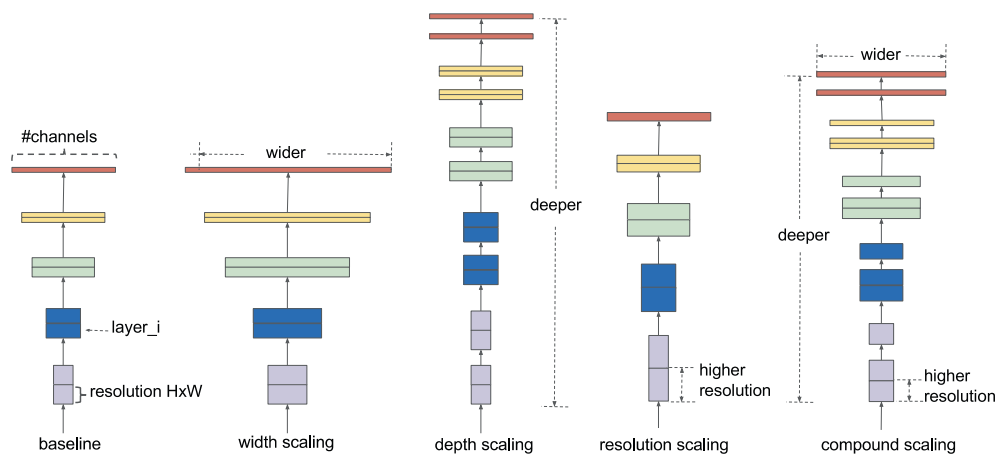
FIGURE 1 Example of scaling types, from left to right: a baseline network example, conventional scaling methods that only increase one network dimension (width, depth, resolution) and, at the end, the EfficientNet compound scaling method that uniformly scales all three dimensions with a fixed ratio. Image taken from the original paper (Tan & Le, 2019).

$$\text{Output} = mode(\text{Output}_1, \text{Output}_2, \dots, \text{Output}_N)$$
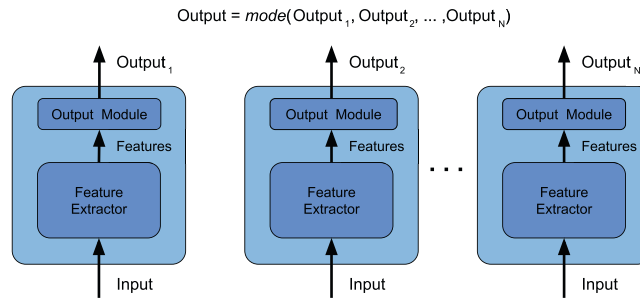
FIGURE 2 Ensemble by voting: the final output is obtained by picking the mode (i.e., most frequent class value) among the output produced by the weak learners. In this way, the weak learners are independent and voting is effective with a high number of heterogeneous weak learners.
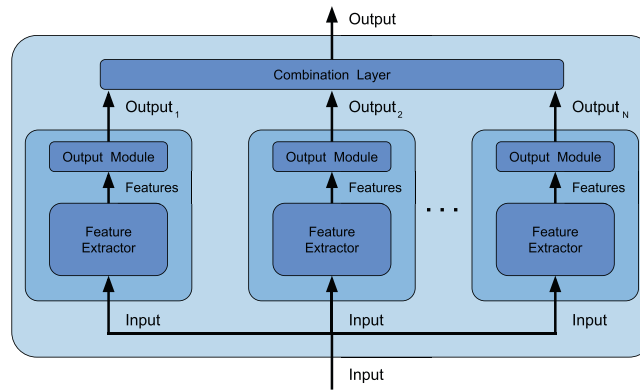
FIGURE 3 Ensemble by output combination: an additional combination layer is fed with the outputs of the weak learners and combines them. In this way, the weak learners are no longer independent and the combination layer can be trained to better adapt to data.

average of them), respectively for classification/regression. However, in this case, the ensemble is only a static aggregator. In this work, we opted for performing an adaptive combination. However, instead of combining the outputs (Figure 3) of the weak learners, we combine the features that the CNNs extract from the input (see Figure 4 where the case $N = 2$ is reported). More formally, let Feat_{weak_i} be the feature vector provided by feature extractor of the i-th weak learner and

$$\text{Feat}_{concat} = \text{Feat}_{weak_1} \oplus \text{Feat}_{weak_2} \oplus \dots \oplus \text{Feat}_{weak_{N-1}} \oplus \text{Feat}_{weak_N} \tag{2}$$

be the vector contained by their concatenation. Then, the final fully connected final layer acts on the combined feature vector Feat_{comb} defined as:

$$\text{Feat}_{comb} = W \cdot \text{Feat}_{concat} + b \tag{3}$$

In this way, we further reduce the complexity of the ensemble without reducing its power and expressiveness. Indeed, the combination layer is of the same type as the output layer of the weak learners (i.e., Linear + LogSoftmax), and keeping both would introduce redundancy. This can be seen as a fully-differentiable version of Gradient Boosting (Friedman, 2000). However, in this way, there is no reason to perform the tree decision traversal, and the ensemble is performed at the features level.

4 | EXPERIMENTAL RESULTS

In this section, the results obtained on several major benchmark datasets on image classification are described. Before showing the results, the main aspects of the experimental setup are detailed. The experiments have been implemented using the PyTorch (Paszke et al., 2019) open-source machine learning framework.
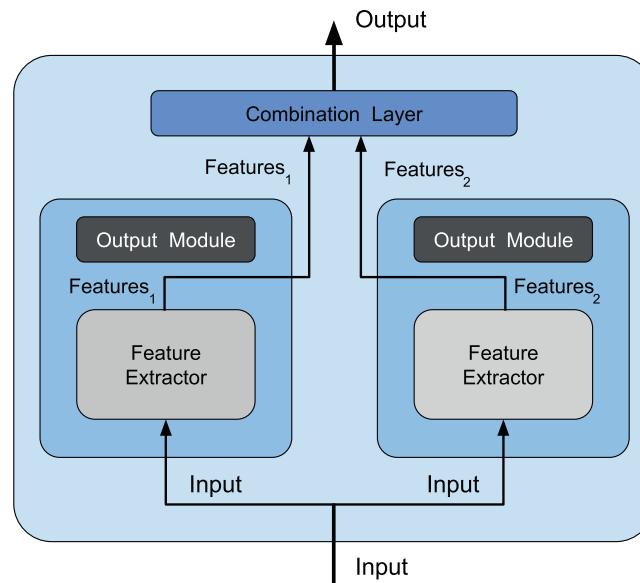
FIGURE 4 Our adaptive ensemble method: is an optimized version of the method shown in Figure 3 because we avoid redundancy and reduce complexity by deleting the output module (dark grey-filled) of weak learners and feeding the combination layer with the features. Light grey-filled modules denote modules whose parameters are frozen during training. The diagram depicts the case $N = 2$, which is used in most of the experiments in this paper, but the method can be applied with an arbitrary value for N.

TABLE 2 Details about the datasets used in the experiments.

Dataset	Domain	Input size	Classes	Balanced	Provided splits
CIFAR-10	Mixed (RGB)	32×32	10	Yes	Train-Test
CIFAR-100	Mixed (RGB)	32×32	100	Yes	Train-Test
Cars	Cars (RGB)	360×240	196	Yes	Train-Test
Food-101	Food (RGB)	512 larger side	101	No	Train-Test
Flower102	Flowers (RGB)	Various	102	Yes	Train-Valid-Test
CINIC-10	Mixed (RGB)	32×32	10	Yes	Train-Valid-Test
Pets	Dogs & Cats (RGB)	Various	37	Yes	Train-Valid-Test

4.1 | Datasets

The proposed solution has been tested on several datasets in order to evaluate its capability of being effective over disparate domains (e.g., type of images, number of classes, balancing, quality) as shown in Table 2. A brief description of each dataset follows.

4.1.1 | CIFAR-10 and CIFAR-100 (Krizhevsky et al., n.d.)

The CIFAR-10 dataset consists of 60,000 32×32 colour images in 10 classes, with 6000 images per class. There are 50,000 training images and 10,000 test images. CIFAR-100 dataset is just like the CIFAR-10, except it has 100 classes containing 600 images each. There are 500 training images and 100 testing images per class. The 100 classes in the CIFAR-100 are grouped into 20 superclasses. Each image comes with a 'fine' label (the class to which it belongs) and a 'coarse' label (the superclass to which it belongs). In the experiments, the fine-grained version with 100 classes has been used.

4.1.2 | Stanford cars (Krause et al., 2013)

The Stanford cars dataset contains 16,185 360×240 colour images of 196 classes of cars at the level of Make, Model, Year (e.g., Tesla, Model S, 2012). The data is split into 8144 training images and 8041 testing images, where each class has been divided roughly in a 50–50 split. Since now, this dataset is referred as 'Cars'.

4.1.3 | Food-101 (Bossard et al., 2014)

The Food-101 dataset consists of 101 food categories with 750 training and 250 test manually-reviewed images per category, making a total of 101,000 images. On purpose, the training images contain some amount of noise that comes mainly in the form of intense colours and sometimes wrong labels. All images were rescaled to have a maximum side length of 512 pixels.

4.1.4 | Oxford 102 flower (Nilsback & Zisserman, 2008)

The Oxford 102 flower is an image classification dataset consisting of 102 flower categories, most of them being plants commonly occurring in the United Kingdom. Each class consists of between 40 and 258 images. The images have large scale, pose and light variations. In addition, there are categories that have significant variations within the category and several very similar ones. Since now, this dataset is referred as 'Flower102'.

4.1.5 | CINIC-10 (Darlow et al., 2018)

CINIC-10 is a dataset for image classification consisting of 270,000 32×32 colour images. It was compiled as a 'bridge' between CIFAR-10 and ImageNet, taking 60,000 images from the former and 210,000 downsampled images from the latter. It is split into three equal subsets—train, validation, and test—each containing 90,000 images.

4.1.6 | Oxford-IIIT pet (Parkhi et al., 2012)

The Oxford-IIIT pet dataset has 37 categories with roughly 200 images for each class representing dogs or cats (25 classes for dogs and 12 for cats). Different versions of the dataset can be used for image classification, object detection, or image segmentation. In particular, for the experimentation, the fine-grained version of the image classification task has been used (i.e., predict the particular breed of the animal in the image instead of just determining if it is a dog or a cat). The images have wide variations in scale, pose and lighting. Since now, this dataset is referred as 'Pets'.

4.2 | Input preprocessing

The models are not fed directly with the images provided by the datasets, but images are preprocessed to improve the performances. In particular, the only two preprocessing steps done are resize (size chosen after preliminary tests) and standardization (in order to have all data of the same dataset described under the same distribution with pixel values centred around the mean and unit deviation) which improves stability and convergence of the training. Preprocessing details for each dataset are shown in Table 3. Even if augmentation has been performed in the works reported as SOTA, no augmentation is performed in this work to test the performances of the 'pure' method.

4.3 | Transfer learning

Transfer learning (Weiss et al., 2016) is the technique of taking knowledge gained while solving one problem, and applying it to a different but related problem. Like most cases for image classification, the stored knowledge is brought by pre-trained models from the ImageNet (Deng

TABLE 3 Input sizes and standardization values, for each channel, used for data preprocessing.

Dataset	Input size	Means (R, G, B)	Stds (R, G, B)
CIFAR-10	256×256	(0.491, 0.482, 0.447)	(0.246, 0.243, 0.261)
CIFAR-100	256×256	(0.507, 0.486, 0.441)	(0.267, 0.256, 0.276)
Cars	500×500	(0.468, 0.457, 0.450)	(0.295, 0.294, 0.302)
Food-101	500×500	(0.550, 0.445, 0.344)	(0.271, 0.275, 0.279)
Flower102	500×500	(0.433, 0.375, 0.285)	(0.296, 0.245, 0.269)
CINIC-10	256×256	(0.478, 0.472, 0.430)	(0.242, 0.238, 0.258)
Pets	500×500	(0.481, 0.448, 0.394)	(0.269, 0.264, 0.272)

et al., 2009) task, since it has more than 14 million images belonging to 1000 generic classes. Transfer learning has been used for weak learners training only.

4.4 | Validation phases

Validation is divided into 2 main phases: end-to-end weak learner overfitting training and ensemble combination layer fine-tuning. In the first phase the transfer learning starts from the ImageNet pre-trained model and sets a new output module (to fit the output size). The models are trained to reach overfit in order to get high specialization on the subset they are referred to. In the second phase, as shown in Figure 4 the weak learners are frozen removing their output modules, so in this phase only the combination layer is trained.

Both phases are performed using the AdaBelief (Zhuang et al., 2020) optimizer which guarantees both fast convergence and generalization. AdaBelief parameters used are the following ones: learning rate 5×10^{-4}, betas (0.9, 0.999), eps 10^{-16}, using weight decoupling without rectify.

4.5 | Avoid overfitting

In order to prevent overfitting (i.e., avoid the model being too specialized to data from the training set with poor performances on unknown data), we use early stopping (i.e., stop training after no improvements on the validation set after a certain number of epochs, called patience) during ensemble fine-tuning only.

4.6 | Data splitting

Every dataset is provided with the 'official' train-test split that is used for the ensemble fine-tuning. On the other hand, for the end-to-end over-fitting training of the weak learners, we perform the following data split:

1. Set the size N of the final ensemble model (i.e., the number of weak learners to be used in the ensembl ing): in particular for the experiments $N = 2$ in order to have the minimum ensemble size;
2. Randomly split the training set into N equally sized and disjoint (i.e., each data belongs exactly only to 1 subset) subsets with stratification (i.e., preserving the class ratios within the subset). During the test only an exception w as made for the Pets dataset, in which the 2 disjoint subsets were made only by cats and dogs, respectively;
3. For each subset, instantiate a weak learner and train it only on that subset (called bagging), with overfitting. In this way every weak learner will be highly specialized only on that portion of data; this could sound self-defeating but (Sollich & Krogh, 1995) has shown that it leads to a qualitative ensemble, especially in the case of this work in which ensembling is adaptive. The choice to reach overfitting will reduce the overall validation time: on the basis of preliminary tests, we noticed that EfficientNet-b0 and AdaBelief optimizer with overfitting training are powerful and will always converge to the same minimum point (very likely to be the global one, due to the fact that accuracy is 100% almost always) independently on the initialization. In this way, just 2 train runs (only one initialization for each weak learner) are sufficient for every dataset.

4.7 | Loss and metrics

4.7.1 | Training loss

Due to the multiclass nature of all dataset tasks, the cross-entropy loss (which exponentially penalizes differences between predicted and true values, expressed as the probability of class belonging) is used. For this reason, the model output has a specified size depending on the dataset (i.e., the number of classes) and each element output[i] represents the probability that the input sample belongs to class i.

4.7.2 | Validation and test metrics

For the validation set evaluation, we decided to use the Weighted F1-score because this takes into account both correct and wrong predictions (true/false positive/negative) and weighting allows to manage any imbalance of the classes (more representative classes have a greater contribution). On the other hand, to make comparisons with previous works on the test set, we used the same metric, which is Accuracy (i.e., correct prediction/total set) in all cases.

4.8 | Hyperparameters

Some hyperparameters have already been fixed and provided in the previous sections (i.e., preprocessing size and standardization, optimizer parameters and ensemble size).To further reduce the total validation time other hyperparameters have been fixed: early stopping patience was set to 10 epochs, batch size to 55 (200 in the case of fine-tuning) and 200 (700 in the case of the fine-tuning) for the 500×500 and 256×256 images, respectively.

Here follows the hyperparameters configuration file for training a weak model, the same format is used for the ensemble, with the only difference that the ensemble_module_list parameter is not empty but contains the local addresses of the two best weak models:

```
project: projects/cifar10 # it varies depending on the dataset
seed: 9999 # it changes for all the runs
# means and standard deviations used for normalization varies depending on the dataset
means: [0.4918687788500817, 0.4826539051649305, 0.44717727749693625]
stds: [0.24697121432552785, 0.24338893940435022, 0.2615925905215076]
early_stopping_patience: 10
num_epochs: 100 # the maximum number of epoch (never reached)
image_size: 256 # size of the images depending on the dataset
batch_size: 200
optim: AdaBelief # the optimizer used
lr: 5e-4 # optimizer parameter
eps: 1e-16 # optimizer parameter
validation_metric: F1 # F1-score is used as validation metric
from_pretrained: True # EfficientNet-b0 pretrained model from ImageNet is used
modeltype: efficientnet-b0
train_ratio: 0.8
valid_ratio: None # automatically obtained
test_ratio: None # automatically obtained
ensemble_module_list: # in case of the ensemble it contains the local addresses of the weak models
```

As written before, there is no hyperparameters tuning, they are all prefixed except for the seeds.

For the ensemble fine-tuning, 5 different random seeds are used. In this way, for each dataset, 2 end-to-end weak training (1 for each subset) and 5 fine-tuning ensemble training are performed.

5 | RESULTS AND DISCUSSION

In this section, the results of the experiments are shown and discussed. Table 4 shows that our work improves the SOTA in all major benchmark datasets and as expected the highest improvements (>0.5%) are obtained on the tasks which are not saturated (i.e., accuracy <99%). These results gain more evidence when complexity is considered too: indeed Table 5 shows that our work (except in the case of CINIC-10) has 5–60 times less total number of parameters and needs 10–100 times fewer FLOPs respect to the SOTA. Moreover, in terms of trainable parameters, since it performs the fine-tuning of a combination layer, our final solution has only about 100 K parameters to train.

In order to stress our method, we also provide a different combination of weak classifiers: specifically, we show the results of an ensemble of five weak models. For demonstration purposes we report the results obtained only for the CIFAR-100 and CIFAR-10 datasets. In the case of CIFAR-100, while the ensemble using 2 weak models obtained an accuracy of 96.808%, the new one obtained an accuracy of 84.930%. This result was expected, since each weak model had to be trained on a third of the images of the previous case according to the data splitting procedure described in Section 4.6 in order to avoid the use of the same images. In the case of CIFAR-10, while the ensemble using 2 weak models obtained an accuracy of 99.612%, the new one obtained an accuracy of 96.640%.

Again both for CIFAR-10 and for CIFAR-100, we also run 6 EfficientNet-b0 weak models and then 6 classical ensembles by majority voting in order to further compare the classical method with ours: one ensemble collects the best two weak models and another the best five ones. For CIFAR-10, the best weak model reaches 97.37% of accuracy, the best ensemble of 2 weak models reaches 97.54% and the ensemble of 5 weak models 97.66% (our method reaches 99.61%). For CIFAR-100 the best weak model reaches 85.55% of accuracy, the ensemble of 2 weak models reaches 86.64% and the best ensemble of 5 weak models 87.56% (our method reaches 96.81%). Moreover, we also run our ensemble method on these classical weak models to show that, as we described in Section 3, our solution improves the results both for the novelties applied to the

TABLE 4 Classification test accuracy comparison between SOTA and our work on datasets used during experiments.

Dataset	SOTA accuracy	Our accuracy	Improvement
CIFAR-10 (Dosovitskiy et al., 2021)	99.500%	99.612%	0.112%
CIFAR-100 (Foret et al., 2021)	96.080%	96.808%	0.728%
Cars (Ridnik et al., 2021)	96.320%	96.868%	0.548%
Food-101 (Foret et al., 2021)	96.180%	96.879%	0.699%
Flower102 (Wu et al., 2021)	99.720%	99.847%	0.127%
CINIC-10 (Lu et al., 2021)	94.300%	95.064%	0.764%
Pets (Foret et al., 2021)	97.100%	98.220%	1.120%

TABLE 5 Complexity, both number of parameters and FLOPs, comparison between SOTA and our work on datasets used during experiments.

Dataset	SOTA parameters	Our parameters	SOTA FLOPs	Our FLOPs
CIFAR-10 (Dosovitskiy et al., 2021)	≈632 M	≈11 M (100 K)	≈916 G[a]	≈0.9 G
CIFAR-100 (Foret et al., 2021)	≈480 M	≈11 M (100 K)	≈299 G[b]	≈0.9 G
Cars (Ridnik et al., 2021)	≈54.7 M	≈11 M (100 K)	≈10 G	≈0.9 G
Food-101 (Foret et al., 2021)	≈480 M	≈11 M (100 K)	≈299 G[b]	≈0.9 G
Flower102 (Wu et al., 2021)	≈277 M	≈11 M (100 K)	≈60 G	≈0.9 G
CINIC-10 (Lu et al., 2021)	≈8.1 M	≈11 M (100 K)	≈1 G	≈0.9 G
Pets (Foret et al., 2021)	≈480 M	≈11 M (100 K)	≈299 G[b]	≈0.9 G

[a]Estimation based on a similar architecture with a similar number of parameters.
[b]Estimation based on the same architecture but scaling FLOPs w.r.t. the number of parameters ratio.

weak models and for those applied to the ensemble. With CIFAR-10 the best adaptive ensemble reaches 97.49% (against our full method that reaches 99.61%) and with CIFAR-100 the best adaptive ensemble reaches 86.79% (against our full method that reaches 96.81%).

Last but not least, below we present an analysis of computation time for a single task; let:

1. T_{end} the time for a single end-to-end weak learner training;
2. T_{fine} the time for a single fine-tuning ensemble model training;
3. T_{fwd} the time for a single forward step;
4. T_{back} the time for a single backward step, when subscripted it indicates the number of parameters involved;
5. T_{upd} the time for a single optimization update step, when subscripted it indicates the number of parameters involved.

Then, for a single task, the total time needed is:

$$T = A \cdot T_{end} + B \cdot T_{fine} \qquad (4)$$

where in our case $A = 2$ since end-to-end training is performed once on each of the two disjoint subsets and $B = 5$ because we performed fine-tuning ensemble training with five random initializations.

However, it is possible to perform in parallel each of the end-to-end training processes, halving the batch size and about taking half of the time; the same goes for the fine-tuning training running all in parallel, in this way the total time is:

$$T = T_{end} + T_{fine} \qquad (5)$$

and considering that a single training is made of forward + backward + update steps to all training data for several epochs:

$$T_{end} \propto T_{fwd} + T_{back} + T_{upd} \qquad (6)$$

$$T_{\text{fine}} \propto N \cdot T_{\text{fwd}} + T_{\text{back}_{100k}} + T_{\text{upd}_{100k}} \underset{\sim}{\propto} T_{\text{fwd}} \tag{7}$$

that is the time for a single fine-tuning ensemble model training is proportional (depending on the actual number of epochs) to the number of the weak learners N multiplied for the time needed for a single forward step, plus the time for a single backward step using 100,000 parameters, plus the time for a single optimization update step using 100,000 parameters that is approximately proportional to the time for a single forward step. Indeed, the approximation in the Equation 7 is justified by the fact that backward and update steps involve only a small fraction of the parameters; moreover, the two weak learners perform forward steps in parallel since they are independent (otherwise we should have $K = 2$). Putting together Equations 6 and 7, the total time is:

$$T \underset{\sim}{\propto} 2 \cdot T_{\text{fwd}} + T_{\text{back}} + T_{\text{upd}} \tag{8}$$

that is, the total time is proportional to 2 multiplied the time for a single forward step plus the time for a single backward step plus the time for a single optimization update step.

What said before, in terms of FLOPs is (considering only one input, just add the linear scaling factor for the training on the whole dataset):

$$F_{\text{fwd}} = F_{\text{back}} = 0.39 \text{ GFLOPs} \tag{9}$$

$$F_{\text{upd}} \approx 20 \cdot P \approx 0.1 \text{ GFLOPs} \tag{10}$$

the Equation 9 refers to FLOPs of EfficientNet-b0 architectures and the Equation 10 refers to FLOPs of AdaBelief update step where $P = 5\,M$ is the number parameters involved in the end-to-end training. Putting all together:

$$\begin{aligned} F &\approx 2 \cdot F_{\text{fwd}} + F_{\text{back}} + F_{\text{upd}} \approx \\ &\approx 2 \cdot 0.39 + 0.39 + 0.1 \approx \\ &\approx 1.3 \text{ GFLOPs} \end{aligned} \tag{11}$$

this means that the whole pipeline on a single image requires about 1.3 GFLOPs, and considering the Table 5, the SOTA for CINIC-10 in (Lu et al., 2021) that has the least number of parameters (8.1 M) requires 1 GFLOPs for one single forward on an image, showing that our solution is the fastest and the speedup is much more noticeable (10–100 times) over the even more complex SOTA models.

6 | CONCLUSION AND FUTURE WORKS

In this work, we presented a method to reverse the trend in image classification of having minor improvements with a huge complexity increase. In particular, we showed a revisited ensembling to outperform the SOTA with restrained complexity, both in terms of the number of parameters and FLOPs. Specifically, we proved how it is possible to perform bagging on two disjoint subsets of data using two EfficientNet-b0 weak learners and training them to overfit on the assigned/scheduled subset.

In this work we pushed the ensemble size to the lower bound using only 2 weak learners: this adaptive ensemble strategy would still be the most efficient using up to 5 weak learners (taking into account that, when using the overfitting strategy, each weak learner has too fragmented and limited knowledge), and then it could be further improved by defining different bagging strategies (e.g., train weak learners on subsets split by class dimensionality, clustering or different colour space mapping of inputs).

Then, the ensemble is performed by fine-tuning a trainable combination layer. The efficiency of the method is given by different reasons: efficiency of EfficientNet-b0 models, fine-tuning for ensemble and the high parallelization capability of the solution, the reduced number of FLOPs combined with the tiny validation space (7 total runs: 2 end-to-end + 5 fine-tuning).

These results pave to investigate this kind of strategy in many fields: Object Detection (performing the ensemble at feature extraction backbone level) and Segmentation (performing the ensemble on the encoding in typical encoder-decoder architectures).

AUTHOR CONTRIBUTIONS

All authors equally contributed to all the sections of this article.

CONFLICT OF INTEREST STATEMENT

The authors declare no conflicts of interest.

DATA AVAILABILITY STATEMENT

ORCID

Bruno Antonio https://orcid.org/0000-0002-9152-8112

Davide Moroni https://orcid.org/0000-0002-5175-5126

Massimo Martinelli https://orcid.org/0000-0001-7419-5099

REFERENCES

Bonawitz, K., Eichner, H., Grieskamp, W., Huba, D., Ingerman, A., Ivanov, V., Kiddon, C., Konečný, J., Mazzocchi, S., McMahan, B., & van Overveldt, T. (2019). Towards federated learning at scale: System design. *Proceedings of Machine Learning and Systems*, 1, 374–388.

Bossard, L., Guillaumin, M., & Van Gool, L. (2014). Food-101 – Mining discriminative components with random forests.

Brock, A., De, S., Smith, S. L., & Simonyan, K. (2021). High-performance large-scale image recognition without normalization. CoRR; abs/2102.06171.

Dai, Z., Liu, H., Le, Q. V., & Tan, M. (2021). CoAtNet: Marrying convolution and attention for all data sizes. CoRR; abs/2106.04803.

Darlow, L. N., Crowley, E. J., Antoniou, A., & Storkey, A. J. (2018). CINIC-10 is not ImageNet or CIFAR-10. ArXiv; abs/1810.03505.

Deng, J., Dong, W., Socher, R., Li, L., Li, K., & Fei-Fei, L. (2009). ImageNet: A large-scale image database: 248–255.

Dhar, P. (2020). The carbon impact of artificial intelligence. *Nature Machine Intelligence*, 2(8), 423–425.

Dong, X., Yu, Z., Cao, W., Shi, Y., & Ma, Q. (2020). A survey on ensemble learning. *Frontiers of Computer Science*, 14(2), 241–258.

Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., & Uszkoreit, J. (2021). An image is worth 16×16 words: Transformers for image recognition at scale.

Foret, P., Kleiner, A., Mobahi, H., & Neyshabur, B. (2021). Sharpness-aware minimization for efficiently improving generalization.

Friedman, J. H. (2000). Greedy function approximation: A gradient boosting machine. *Annals of Statistics*, 29, 1189–1232.

Ganaie, M. A., Hu, M., Malik, A. K., Tanveer, M., & Suganthan, P. N. (2021). Ensemble deep learning: A review. arXiv preprint arXiv:2104.02395.

Goel, A., Tung, C., Lu, Y. H., & Thiruvathukal, G. K. (2020). A survey of methods for low-power deep learning and computer vision. IEEE: 1–6.

Krause, J., Stark, M., Deng, J., & Fei-Fei, L. (2013). 3D object representations for fine-grained categorization.

Krizhevsky, A., Nair, V., & Hinton, G. CIFAR-10 (Canadian Institute for Advanced Research).

Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). ImageNet classification with deep convolutional neural networks. *Communications of the ACM*, 60, 84–90.

Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., & Guo, B. (2021). Swin transformer: Hierarchical vision transformer using shifted windows. CoRR; abs/2103.14030.

Lu, Z., Sreekumar, G., Goodman, E., Banzhaf, W., Deb, K., & Boddeti, V. N. (2021). Neural architecture transfer. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(9), 2971–2989. https://doi.org/10.1109/tpami.2021.3052758

Nilsback, M. E., & Zisserman, A. (2008). Automated flower classification over a large number of classes.

Opitz, D. W., & Maclin, R. (1999). Popular ensemble methods: An empirical study. *Journal of Artificial Intelligence Research*, 11, 169–198. https://doi.org/10.1613/jair.614

Parkhi, O. M., Vedaldi, A., Zisserman, A., & Jawahar, C. V. (2012). Cats and dogs.

Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., & Desmaison, A. (2019). PyTorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, D. F. Alché-Buc, E. Fox, & R. Garnett (Eds.), *Advances in neural information processing systems* (Vol. 32, pp. 8024–8035). Curran Associates, Inc.

Ridnik, T., Ben-Baruch, E., Noy, A., & Zelnik-Manor, L. (2021). ImageNet-21K pretraining for the masses.

Sagi, O., & Rokach, L. (2018). Ensemble learning: A survey. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 8, e1249. https://doi.org/10.1002/widm.1249

Schwartz, R., Dodge, J., Smith, N. A., & Etzioni, O. (2020). Green ai. *Communications of the ACM*, 63(12), 54–63.

Sollich, P., & Krogh, A. (1995). Learning with ensembles: How over-fitting can be useful. In *NIPS'95* (pp. 190–196). MIT Press.

Sze, V., Chen, Y. H., Emer, J., Suleiman, A., & Zhang, Z. (2017). Hardware for machine learning: Challenges and opportunities. IEEE: 1–8.

Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., & Wojna, Z. (2016). Rethinking the inception architecture for computer vision: 2818–2826.

Tan, M., & Le, Q. (2019). EfficientNet: Rethinking model scaling for convolutional neural networks. In K. Chaudhuri & R. Salakhutdinov (Eds.), *Proceedings of the 36th international conference on machine learning. 97 of proceedings of machine learning research* (pp. 6105–6114). PMLR.

Weiss, K., Khoshgoftaar, T., & Wang, D. (2016). A survey of transfer learning. *Journal of Big Data*, 3, 1–40. https://doi.org/10.1186/s40537-016-0043-6

Wu, H., Xiao, B., Codella, N., Liu, M., Dai, X., Yuan, L., & Zhang, L. (2021). CvT: Introducing convolutions to vision transformers.

Xie, S., Girshick, R., Dollár, P., Tu, Z., & He, K. (2017). Aggregated residual transformations for deep neural networks: 5987–5995.

Zhai, X., Kolesnikov, A., Houlsby, N., & Beyer, L. (2021). Scaling vision transformers. ArXiv; abs/2106.04560.

Zhuang, J., Tang, T., Ding, Y., Tatikonda, S. C., Dvornek, N., Papademetris, X., & Duncan, J. (2020). AdaBelief optimizer: Adapting stepsizes by the belief in observed gradients. *Conference on Neural Information Processing Systems*, 33, 18795–18806.

AUTHOR BIOGRAPHIES

Antonio Bruno. He received the Master degree in Computer Science from the University of Pisa. He received a research grant from the Signal and Images Lab (SI-Lab) of the Institute of Information Science and Technologies (ISTI) for collaborating on the Barilla AGROSAT Plus project. His main research interests include deep learning for structured domain (e.g., sequences, trees, images).

Davide Moroni. He received the MSc degree (Hons.) in mathematics from the University of Pisa, in 2001, the Diploma from the Scuola Normale Superiore of Pisa, in 2002, and the PhD degree in mathematics from the University of Rome La Sapienza, in 2006. He is a researcher with the Institute of Information Science and Technologies (ISTI), National Research Council, Italy, Pisa. He is currently the head of the Signals and Images Lab, ISTI. He is the chair of the MUSCLE working group of the European Consortium for Informatics and Mathematics. Since 2018, he serves as the chair of the Technical Committee 16 on Algebraic and Discrete Mathematical Techniques in Pattern Recognition and Image Analysis of the International Association for Pattern Recognition (IAPR). He is an associate editor of IET Image Processing. His main research interests include geometric modelling, computational topology, image processing, computer vision, and medical imaging. At the moment, he is leading the ISTI-CNR team in the National Project PON MIUR S4E, working on maritime safety and security, and in the regional Project IRIDE addressing AR technologies and computer vision of Industry 4.0.

Massimo Martinelli. He is member of the Signals & Images research laboratory at the Institute of Information Science and Technologies (ISTI), National Research Council (CNR), Italy, Pisa, since 1987. Head of the 'Artificial Intelligence Technologies and Frameworks Area' at SI-Lab since 2017. He was member of the W3C Multimedia Semantics Incubator Group (2006–2007). He is currently leading the CNR-ISTI team in the projects TiAssisto (Tuscany Region), Barilla Agrosat Plus (industrial), Cloud Pathology (industrial), and of the Scientific Agreements with the UO Otolaryngology, audiology and phoniatrics (UNIPI), and with the Italian Mountain Medicine Society. Member of the Doseteam4you group of the Department of Diagnostics and Interventional Radiology of the University Hospital of Pisa. Topic Editor of 'Machine Learning and Biomedical Sensors' of the Sensors Journal, His main scientific interests include Computer Vision, Deep Learning, Decision Support Systems, Web technologies.