

Boosted Residual Networks

Alan Mosca^(✉) and George D. Magoulas

Department of Computer Science and Information Systems,
Birkbeck, University of London, London, UK
{a.mosca,gmagoulas}@dcs.bb.k.ac.uk

Abstract. In this paper we present a new ensemble method, called Boosted Residual Networks, which builds an ensemble of Residual Networks by growing the member network at each round of boosting. The proposed approach combines recent developments in Residual Networks - a method for creating very deep networks by including a shortcut layer between different groups of layers - with the Deep Incremental Boosting, which has been proposed as a methodology to train fast ensembles of networks of increasing depth through the use of boosting. We demonstrate that the synergy of Residual Networks and Deep Incremental Boosting has better potential than simply boosting a Residual Network of fixed structure or using the equivalent Deep Incremental Boosting without the shortcut layers.

1 Introduction

Residual Networks, a type of deep network recently introduced in [3], are characterized by the use of *shortcut* connections (sometimes also called *skip* connections), which connect the input of a layer of a deep network to the output of another layer positioned a number of levels “above” it. The result is that each one of these shortcuts shows that networks can be build in *blocks*, which rely on both the output of the previous layer and the previous block.

Residual Networks have been developed with many more layers than traditional Deep Networks, in some cases with over 1000 blocks, such as the networks in [5]. A recent study in [14] compares Residual Networks to an ensemble of smaller networks. This is done by unfolding the shortcut connections into the equivalent tree structure, which closely resembles an ensemble. An example of this can be shown in Fig. 1.

Dense Convolutional Neural Networks [6] are another type of network that makes use of shortcuts, with the difference that each layer is connected to all its ancestor layers directly by a shortcut. Similarly, these could be also unfolded into an equivalent ensemble.

True ensemble methods are often left as an *afterthought* in Deep Learning models: it is generally considered sufficient to treat the Deep Learning method

G.D. Magoulas—The authors gratefully acknowledge the support of NVIDIA Corporation with the donation of the Tesla Titan X Pascal GPUs used for this research.

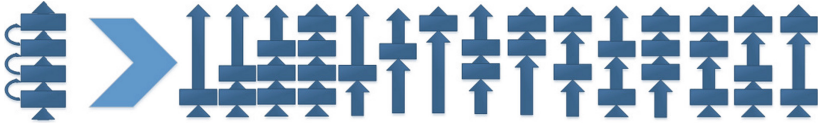


Fig. 1. A Residual Network of N blocks can be unfolded into an ensemble of $2^N - 1$ smaller networks.

as a “black-box” and use a well-known generic Ensemble method to obtain marginal improvements on the original results. Whilst this is an effective way of improving on existing results without much additional effort, we find that it can amount to a waste of computations. Instead, it would be much better to apply an Ensemble method that is aware, and makes use of, the underlying Deep Learning algorithm’s architecture.

We define such methods as “white-box” Ensembles, which allow us to improve on the generalisation and training speed compared to traditional Ensembles, by making use of particular properties of the base classifier’s learning algorithm and architecture. We propose a new such method, which we call Boosted Residual Networks (BRN), which makes use of developments in Deep Learning, previous other white-box Ensembles and combines several ideas to achieve improved results on benchmark datasets.

Using a white-box ensemble allows us to improve on the generalisation and training speed of other ensemble methods by making use of the knowledge of the base classifier’s structure and architecture. Experimental results show that Boosted Residual Networks achieves improved results on benchmark datasets.

The next section presents the background on Deep Incremental Boosting. Then the proposed Boosted Residual Networks method is described. Experiments and results are discussed next, and the paper ends with conclusions.

2 Background

Deep Incremental Boosting, introduced in [9], is an example of such white-box ensemble method developed for building ensembles Convolutional Networks. The method makes use of principles from transfer of learning, like for example those used in [15], applying them to conventional AdaBoost [12]. Deep Incremental Boosting increases the size of the network at each round by adding new layers at the end of the network, allowing subsequent rounds of boosting to run much faster. In the original paper on Deep Incremental Boosting [9], this has been shown to be an effective way to learn the *corrections* introduced by the emphasisation of learning mistakes of the boosting process. The argument as to why this works effectively is based on the fact that the datasets at rounds t and $t + 1$ will be *mostly similar*, and therefore a classifier h_t that performs better than randomly on the resampled dataset \mathbf{X}_t will also perform better than randomly on the resampled dataset \mathbf{X}_{t+1} . This is under the assumption that both

datasets are sampled from a common ancestor set \mathbf{X}_a . It is subsequently shown that such a classifier can be re-trained on the differences between \mathbf{X}_t and \mathbf{X}_{t+1} .

This practically enables the ensemble algorithm to train the subsequent rounds for a considerably smaller number of epochs, consequently reducing the overall training time by a large factor. The original paper also provides a conjecture-based justification for why it makes sense to extend the previously trained network to learn the ‘‘corrections’’ taught by the boosting algorithm. A high level description of the method is shown in Algorithm 1, and the structure of the network at each round is illustrated in Fig. 2.

Algorithm 1. Deep Incremental Boosting

```

 $D_{0,i} = 1/M$  for all  $i$ 
 $t = 0$ 
 $\mathbf{W}_0 \leftarrow$  randomly initialised weights for first classifier
while  $t < t_{end}$  do
     $\mathbf{X}_t \leftarrow$  sample from  $\mathbf{X}_0$  with distribution  $\mathbf{D}_t$ 
     $u_t \leftarrow$  create untrained classifier with additional layer of shape  $L_{new}$ 
    copy weights from  $\mathbf{W}_t$  into the bottom layers of  $u_t$ 
     $h_t \leftarrow$  train  $u_t$  classifier on current subset
     $\mathbf{W}_{t+1} \leftarrow$  all weights from  $h_t$ 
     $\epsilon_t = \frac{1}{2} \sum_{(i,y) \in B} D_{t,i} (1 - h_t(x_i, y_i) + h_t(x_i, y))$ 
     $\beta_t = \epsilon_t / (1 - \epsilon_t)$ 
     $D_{t+1,i} = \frac{D_{t,i}}{Z_t} \cdot \beta^{(1/2)(1+h_t(x_i, y_i) - h_t(x_i, y))} \forall i$ 
    where  $Z_t$  is a normalisation factor such that  $\mathbf{D}_{t+1}$  is a distribution
     $\alpha_t = \frac{1}{\beta_t}$ 
     $t = t + 1$ 
end while
 $H(x) = \operatorname{argmax}_{y \in Y} \sum_{t=1}^T \log \alpha_t h_t(x, y)$ 
    
```

3 Creating the Boosted Residual Network

In this section we propose a method for generating Boosted Residual Networks. This works by increasing the size of an original residual network by one residual block at each round of boosting. The method achieves this by selecting an *injection point* index p_i at which the new block is to be added, which is not necessarily the last block in the network, and by transferring the weights from the layers below p_i in the network trained at the previous round of boosting.

The boosting method performs an iterative re-weighting of the training set, which skews the resample at each round to *emphasize* the training examples that are harder to train. Therefore, it becomes necessary to utilise the entire ensemble at test time, rather than just use the network trained in the last round. This has the effect that the Boosted Residual Networks cannot be used as a way to train a single Residual Network incrementally. However, as we will discuss

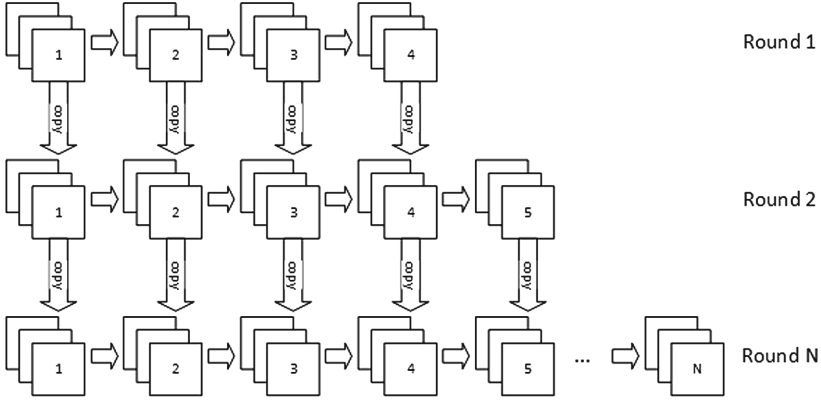


Fig. 2. Illustration of subsequent rounds of Deep Incremental Boosting

later, it is possible to alleviate this situation by deriving an approach that uses bagging instead of boosting; therefore removing the necessity to use the entire ensemble at test time. It is also possible to delete individual blocks from a Residual Network at training and/or testing time, as presented in [3], however this issue is considered out of the scope of this paper.

The iterative algorithm used in the paper is shown in Algorithm 2. At the first round, the entire training set is used to train a network of the original *base* architecture, for a number of epochs n_0 . After the first round, the following steps are taken at each subsequent round t :

- The ensemble constructed so far is evaluated on the training set to obtain the set errors ϵ , so that a new training set can be sampled from the original training set. This is a step common to all boosting algorithms.
- A new network is created, with the addition of a new block of layers B_{new} immediately after position p_t , which is determined as an initial pre-determined position p_0 plus an offset $t * \delta_p$ for all the blocks added at previous layers, where δ_p is generally chosen to be the size of the newly added layers. This puts the new block of layers immediately after the block of layers added at the previous round, so that all new blocks are effectively added sequentially.
- The weights from the layers below p_t are copied from the network trained at round $t - 1$ to the new network. This step allows to considerably shorten the training thanks to the transfer of learning shown in [15].
- The newly created network is subsequently trained for a reduced number of epochs $n_{t>0}$.
- The new network is added to the ensemble following the conventional rules and weight $\alpha_t = \frac{1}{\beta_t}$ used in AdaBoost. We did not see a need to modify the way β_t is calculated, as it has been performing well in both DIB and many AdaBoost variants [2, 9, 12, 13].

Algorithm 2. Boosted Residual Networks

$D_{0,i} = 1/M$ for all i
 $t = 0$
 $W_0 \leftarrow$ randomly initialised weights for first classifier
 $p_0 \leftarrow$ initial injection position
while $t < T$ **do**
 $X_t \leftarrow$ sample from X_0 with distribution D_t
 $u_t \leftarrow$ create untrained classifier with an additional block B_{new} of pre-determined shape N_{new}
 determine block injection position $p_t = p_{t-1} + |B_{new}|$
 connect the input of B_{new} to the output of layer $p_t - 1$
 connect the output of B_{new} and of layer $p_t - 1$ to a merge layer m_i
 connect the merge layer to the remainder of the network
 copy weights from W_t into the bottom layers $l < p_t$ of u_t
 $h_t \leftarrow$ train u_t classifier on current subset
 $W_{t+1} \leftarrow$ all weights from h_t
 $\epsilon_t = \frac{1}{2} \sum_{(i,y) \in B} D_{t,i} (1 - h_t(x_i, y) + h_t(x_i, y))$
 $\beta_t = \epsilon_t / (1 - \epsilon_t)$
 $D_{t+1,i} = \frac{D_{t,i}}{Z_t} \cdot \beta^{(1/2)(1+h_t(x_i,y_i)-h_t(x_i,y))} \forall i$
 where Z_t is a normalisation factor such that D_{t+1} is a distribution
 $\alpha_t = \frac{1}{\beta_t}$
 $t = t + 1$
end while
 $H(x) = \operatorname{argmax}_{y \in Y} \sum_{t=1}^T \log \alpha_t h_t(x, y)$

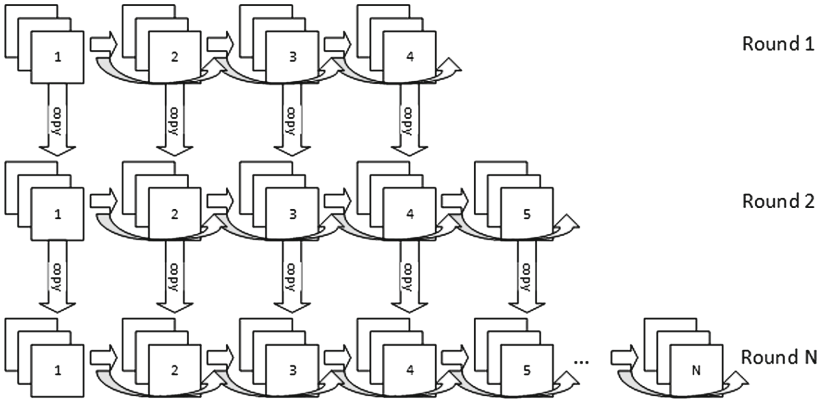


Fig. 3. Illustration of subsequent rounds of Boosted Residual Networks

Figure 3 shows a diagram of how the Ensemble is constructed by deriving the next network at each round of boosting from the network used in the previous round.

We identified a number of optional variations to the algorithm that may be implemented in practice, which we have empirically established as not having

an impact on the overall performance of the network. We report them here for completeness.

- Freezing the layers that have been copied from the previous round.
- Only utilising the weights distribution for the examples in the training set instead of resampling, as an input to the training algorithm.
- Inserting the new block always at the same position, rather than after the previously-inserted block (we found this to affect performance negatively).

3.1 Comparison to Approximate Ensembles

While both Residual Networks and Densely Connected Convolutional Networks may be unfolded into an equivalent ensemble, we note that there is a differentiation between an actual ensemble method and an ensemble “approximation”. During the creation of an ensemble, one of the principal factors is the creation of *diversity*: each base learner is trained independently, on variations (resamples in the case of boosting algorithms) of the training set, so that each classifier is guaranteed to learn a different function that represents an approximation of the training data. This is the enabling factor for the ensemble to perform better in aggregate.

In the case of Densely Connected Convolutional Networks (DCCN) specifically, one may argue that a partial unfolding of the network could be, from a schematic point of view, very similar to an ensemble of incrementally constructed Residual Networks. We make the observation that, although this would be correct, on top of the benefit of diversity, our method also provides a much faster training methodology: the only network that is trained for a full schedule is the network created at the first round, which is also the smallest one. All subsequent networks are trained for a much shorter schedule, saving a considerable amount of time. Additionally, while the schematic may seem identical, there is a subtle difference: each member network outputs a classification of its own, which is then aggregated by a weighted averaging determined by the errors on the test set, whilst in a DCCN the input of the final aggregation layer is the output of each underlying set of layers. We conjecture that this aggressive dimensionality reduction before the aggregation has a regularising effect on the ensemble.

4 Experiments and Discussion

In the experiments we used the MNIST, CIFAR-10 and CIFAR-100 datasets, and compared Boosted Residual Networks (BRN) with an equivalent Deep Incremental Boosting without the skip-connections (DIB), AdaBoost with the initial base network as the base classifier (AdaBoost) and the single Residual Network equivalent to the last round of Boosted Residual Networks (ResNet). It is to be noted that, because of the shortened training schedule and the differing architecture, the results on the augmented datasets are not the same as those reported in the original paper for Residual Networks.

Table 1. Test accuracy in the benchmark datasets for the methods compared.

	ResNet	AdaBoost	DIB	BRN
MNIST	99.41%	99.41%	99.47%	99.53%
CIFAR-10	89.12%	89.74%	90.83%	90.85%
CIFAR-10 (aug)	92.14%	91.66%	92.31%	92.94%
CIFAR-100	67.25%	68.18%	68.56%	69.04%
CIFAR-100 (aug)	69.72%	70.74%	71.55%	72.41%

MNIST [8] is a common computer vision dataset that associates 70000 pre-processed images of hand-written numerical digits with a class label representing that digit. The input features are the raw pixel values for the 28×28 images, in grayscale, and the outputs are the numerical value between 0 and 9. 50000 samples are used for training, 10000 for validation, and 10000 for testing.

CIFAR-10 is a dataset that contains 60000 small images of 10 categories of objects. It was first introduced in [7]. The images are 32×32 pixels, in RGB format. The output categories are *airplane, automobile, bird, cat, deer, dog, frog, horse, ship, truck*. The classes are completely mutually exclusive so that it is translatable to a *1-vs-all* multiclass classification. 50000 samples are used for training, and 10000 for testing. This dataset was originally constructed without a validation set.

CIFAR-100 is a dataset that contains 60000 small images of 100 categories of objects, grouped in 20 super-classes. It was first introduced in [7]. The image format is the same as CIFAR-10. Class labels are provided for the 100 classes as well as the 20 super-classes. A super-class is a category that includes 5 of the fine-grained class labels (e.g. “insects” contains *bee, beetle, butterfly, caterpillar, cockroach*). 50000 samples are used for training, and 10000 for testing. This dataset was originally constructed without a validation set.

In order to reduce noise in the results, we aligned the random initialisation of all networks across experiments, by fixing the seeds for the random number generators. All experiments were repeated 10 times and we report the mean performance values. We also report results with light data augmentation: we randomly rotated, flipped horizontally and scaled images, but did not use any heavy augmentation, including random crops. Results are reported in Table 1, while Fig. 4 shows a side-by-side comparison of accuracy levels at each round of boosting for both DIB and BRN on the MNIST and CIFAR-100 test sets. This figure illustrates how BRNs are able to consistently outperform DIB at each intermediate value of ensemble size, and although such differences would still fall within a Bernoulli confidence interval of 95%, we make the note that this does not take account of the fact that all the random initialisations were aligned, so both methods started with the exact same network. In fact, an additional Friedman Aligned Ranks test on the entire group of algorithms tested shows that there is a statistically significant difference in generalisation performance,

Table 2. Training times comparison. BRN and DIB are the fastest Ensemble methods compared. The time to train the base network and a ResNet of comparable performance is reported for comparison.

	ResNet	Base Net	AdaBoost	DIB	BRN
MNIST	217 min	62 min	442 min	202 min	199 min
CIFAR-10	1941 min	184 min	1212 min	461 min	449 min
CIFAR-10 (aug)	2228 min	213 min	2150 min	1031 min	911 min
CIFAR-100	2172 min	303 min	2873 min	607 min	648 min
CIFAR-100 (aug)	2421 min	328 min	3072 min	751 min	735 min

whilst a direct Wilcoxon test comparing only BRN and DIB shows that BRN is significantly better.

Table 2 shows that this is achieved without significant changes in the training time¹. The main speed increase is due to the fact that the only network being trained with a full schedule is the first network, which is also the smallest, whilst all other derived networks are trained for a much shorter schedule (in this case only 10% of the original training schedule). If we exclude the single network, which is clearly from a different distribution and only mentioned for reference, a Friedman Aligned Ranks test shows that there is a statistically significant difference in speed between the members of the group, but a Wilcoxon test between Deep Incremental Boosting and Boosted Residual Networks does not show a significant difference. This confirms what could be conjured from the algorithm itself for BRN, which is of the same complexity w.r.t. the number of Ensemble members as DIB.

The initial network architectures for the first round of boosting are shown in Table 3a for MNIST, and Table 3b for CIFAR-10 and CIFAR-100. The single networks currently used to reach state-of-the-art on these datasets are very cumbersome in terms of resources and training time. Instead, we used relatively simpler network architectures that were faster to train, which still perform well on the datasets at hand, with accuracy close to, and almost comparable to, the state-of-the-art. This enabled us to test larger Ensembles within an acceptable training time. Our intention is to demonstrate a methodology that makes it feasible to create ensembles of Residual Networks following a “white-box” approach to significantly improve the training times and accuracy levels achievable with current ensemble methods.

Training used the WAME method [11], which has been shown to be faster than Adam and RMSprop, whilst still achieving comparable generalisation. This is thanks to a specific weight-wise learning rate acceleration factor that is determined based only on the sign of the current and previous partial derivative $\frac{\partial E(x)}{\partial w_{i,j}}$. For the single Residual Network, and for the networks in AdaBoost, we trained

¹ In some cases BRN is actually faster than DIB, but we believe this to be just noise due to external factors such as system load.

Table 3. Network structures used in experiments. The layers marked with “*” indicate the location after which we added the residual blocks.

64 conv, 5×5	2×96 conv, 3×3
2×2 max-pooling	96 conv, 3×3 , 2×2 strides
128 conv, 5×5	96 conv, 3×3 , 2×2 strides
2×2 max-pooling *	96 conv, 3×3 , 2×2 strides
Dense, 1024 nodes	2×2 max-pooling
50% dropout	2×192 conv, 3×3
(a) MNIST	192 conv, 3×3 , 2×2 strides
	192 conv, 3×3 , 2×2 strides
	192 conv, 3×3 , 2×2 strides
	2×2 max-pooling *
	192 conv, 3×3
	192 conv, 1×1
	10 conv, 1×1
	global average pooling
	10-way softmax
	(b) CIFAR-10 and CIFAR-100

Table 4. Structure of blocks added at each round of DIB and BRN.

64 conv, 3×3	192 conv, 3×3
Batch Normalization	Batch Normalization
ReLU activation	ReLU activation
(a) MNIST	192 conv, 3×3
	Batch Normalization
	ReLU activation
	(b) CIFAR-10 and CIFAR-100

each member for 100 epochs. For Deep Incremental Boosting and Boosted Residual Networks, we trained the first round for 50 epochs, and every subsequent round for 10 epochs, and ran all the algorithms for 10 rounds of boosting, except for the single network. The structure of each incremental block added to Deep Incremental Boosting and Boosted Residual Networks at each round is shown in Table 4a for MNIST, and in Table 4b for CIFAR-10 and CIFAR-100. All layers were initialised following the recommendations in [4].

Distilled Boosted Residual Network: DBRN In another set of experiments we tested the performance of a Distilled Boosted Residual Network (DBRN). Distillation has been shown to be an effective process for regularising large Ensembles of Convolutional Networks in [10], and we have applied the same methodology to the proposed Boosted Residual Network. For the distilled network structure we used the same architecture as that of the Residual Network from the final round of boosting. Average accuracy results in testing over 10 runs are presented

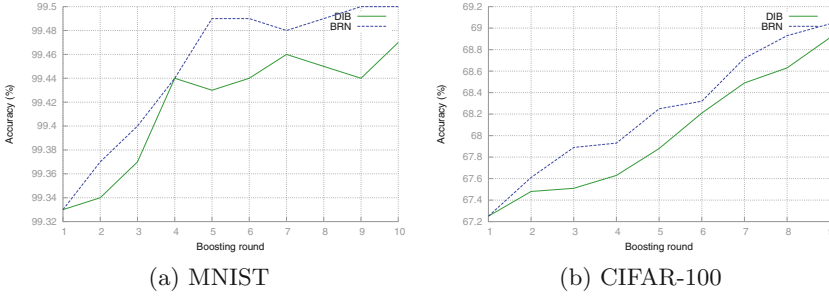


Fig. 4. Round-by-round comparison of DIB vs BRN on the test set

Table 5. Comparative results in terms of testing accuracy.

	DBRN	DDIB
MNIST	99.49%	99.44%
CIFAR-10	91.11%	90.66%
CIFAR-10 (aug)	93.06%	92.71%
CIFAR-100	66.63%	65.91%
CIFAR-100 (aug)	70.24%	69.18%

in Table 5, and for completeness of comparison we also report the results for the distillation of DIB, following the same procedure, as DDIB. DBRN does appear to improve results only for CIFAR-10, but it consistently beats DDIB on all datasets. These differences are too small to be deemed statistically significant, confirming that the function learned by both BRN and DIB can be efficiently transferred to a single network.

Bagged Residual Networks: BARN We experimented with substituting the boosting algorithm with a simpler bagging algorithm [1] to evaluate whether it would be possible to only use the network from the final round of bagging as an approximation of the Ensemble. We called this the Bagged Approximate Residual Networks (BARN) method. We then also tested the performance of the Distilled version of the whole Bagged Ensemble for comparison. The results are reported as “DBARN”. The results are reported in Table 6. It is not clear whether using the last round of bagging is significantly comparable to using the entire Bagging ensemble at test time, or deriving a new distilled network from it, and more experimentation would be required.

Table 6. Test accuracy for BARN and DBARN.

	BRN	Bagging	BARN	DBARN
MNIST	99.53%	99.55%	99.29%	99.36%
CIFAR-10	90.85%	91.43%	88.47%	90.63%
CIFAR-10 (aug)	92.94%	92.61%	92.73%	92.69%
CIFAR-100	69.04%	68.15%	69.42%	66.16%
CIFAR-100 (aug)	72.41%	71.90%	72.01%	70.44%

5 Conclusions and Future Work

In this paper we have derived a new ensemble algorithm specifically tailored to Convolutional Networks to generate Boosted Residual Networks. We have shown that this surpasses the performance of a single Residual Network equivalent to the one trained at the last round of boosting, of an ensemble of such networks trained with AdaBoost, and Deep Incremental Boosting on the MNIST and CIFAR datasets, with and without using augmentation techniques.

We then derived and looked at a distilled version of the method, and how this can serve as an effective way to reduce the test-time cost of running the Ensemble. We used Bagging as a proxy to test the generation of the approximate Residual Network, which, with the parameters tested, does not perform as well as the original Residual Network, BRN or DBRN.

Further experimentation of the Distilled methods presented in the paper, namely DBRN and DBARN, is necessary to fully investigate their behaviour. This is indeed part of our work in the near future. Additionally, the Residual Networks built in our experiments were comparatively smaller than those that achieve state-of-the-art performance. Nevertheless, it might be appealing in the future to evaluate the performance improvements obtained when creating ensembles of larger, state-of-the-art, networks. Additional further investigation could also be conducted on the creation of Boosted Densely Connected Convolutional Networks, by applying the same principle to DCCN instead of Residual Networks.

References

1. Breiman, L.: Bagging predictors. *Mach. Learn.* **24**(2), 123–140 (1996)
2. Freund, Y., Iyer, R., Schapire, R.E., Singer, Y.: An efficient boosting algorithm for combining preferences. *J. Mach. Learn. Res.* **4**, 933–969 (2003)
3. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. arXiv preprint [arXiv:1512.03385](https://arxiv.org/abs/1512.03385) (2015)
4. He, K., Zhang, X., Ren, S., Sun, J.: Delving deep into rectifiers: surpassing human-level performance on imagenet classification. In: *Proceedings of the IEEE International Conference on Computer Vision*, pp. 1026–1034 (2015)
5. He, K., Zhang, X., Ren, S., Sun, J.: Identity mappings in deep residual networks. arXiv preprint [arXiv:1603.05027](https://arxiv.org/abs/1603.05027) (2016)

6. Huang, G., Liu, Z., Weinberger, K.Q.: Densely connected convolutional networks. arXiv preprint [arXiv:1608.06993](https://arxiv.org/abs/1608.06993) (2016)
7. Krizhevsky, A., Hinton, G.: Learning multiple layers of features from tiny images (2009)
8. Lecun, Y., Cortes, C.: The MNIST database of handwritten digits. <http://yann.lecun.com/exdb/mnist/>
9. Mosca, A., Magoulas, G.: Deep incremental boosting. In: Benzmulder, C., Sutcliffe, G., Rojas, R. (eds.) GCAI 2016, 2nd Global Conference on Artificial Intelligence. EPIc Series in Computing, vol. 41, pp. 293–302. EasyChair (2016)
10. Mosca, A., Magoulas, G.D.: Regularizing deep learning ensembles by distillation. In: 6th International Workshop on Combinations of Intelligent Methods and Applications (CIMA 2016), p. 53 (2016)
11. Mosca, A., Magoulas, G.D.: Training convolutional networks with weight-wise adaptive learning rates. In: ESANN 2017 Proceedings, European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning, Bruges, Belgium, 26–28 April 2017 (2017, in press). [i6doc.com](http://www.i6doc.com)
12. Schapire, R.E.: The strength of weak learnability. *Mach. Learn.* **5**, 197–227 (1990)
13. Schapire, R.E., Freund, Y.: Experiments with a new boosting algorithm. In: *Machine Learning: Proceedings of the Thirteenth International Conference*, pp. 148–156 (1996)
14. Veit, A., Wilber, M., Belongie, S.: Residual networks behave like ensembles of relatively shallow networks. arXiv e-prints, May 2016
15. Yosinski, J., Clune, J., Bengio, Y., Lipson, H.: How transferable are features in deep neural networks? In: *Advances in Neural Information Processing Systems*, pp. 3320–3328 (2014)