

Incremental Training of Deep Convolutional Neural Networks

R. Istrate^{1,2}, A. C. I. Malossi¹, C. Bekas¹, and D. Nikolopoulos²

¹ IBM Research – Zurich, Switzerland
{roi,acm,bek}@zurich.ibm.com

² Queen’s University of Belfast, United Kingdom
{ristrate01,d.nikolopoulos}@qub.ac.uk

Abstract. We propose an incremental training method that partitions the original network into sub-networks, which are then gradually incorporated in the running network during the training process. To allow for a smooth dynamic growth of the network, we introduce a look-ahead initialization that outperforms the random initialization. We demonstrate that our incremental approach reaches the reference network baseline accuracy. Additionally, it allows to identify smaller partitions of the original state-of-the-art network, that deliver the same final accuracy, by using only a fraction of the global number of parameters. This allows for a potential speedup of the training time of several factors. We report training results on CIFAR-10 for ResNet and VGGNet.

Keywords: Training algorithm, Look-ahead, CNNs

1 Introduction

When dealing with a classification task on a new dataset of images, a widely used strategy is to start by training a few state-of-the-art networks that were developed for similar datasets. The main disadvantage of this approach is that only too late in the process we learn whether the network is not well suited for the dataset. At that moment, the training is stopped, the network is adapted and the process is restarted from scratch, discarding all the previously collected information. Thus, in the presence of large, complex datasets, the global time-to-solution tends to reach the order of several weeks or even months.

A great deal of research is now focused in optimizing the depth of a neural network (NN) and most of the proposed methods involve growing the network by gradually adding one or more layers [1–3]. Although this is a logical step towards evolving network architectures, to the best of our knowledge there is no work that compares the accuracy and performance of the networks obtained by dynamically adapting their structure during training, with the original one trained from scratch.

In this paper we present our contribution towards gradually training deep state-of-the-art convolutional neural networks (CNNs) with no loss in accuracy. During incremental training we start the learning process with a shallow network.

When the network performance stops improving because of its limited capacity, we transfer the knowledge of the shallow network to a deeper one and continue the training. In this way the time and resources spent in training the initial shallow network are not wasted and the deeper network has a better initialization point. Additionally, since the incremental training method easily quantifies (by construction) the value added by each network extension, it is simple to define custom trade-off functions (e.g., time vs accuracy) to stop the network expansion when desired criteria are met. The long-term goal is to develop a fully-automated framework that optimizes neural network structures for specific tasks. However, this requires knowing what type of changes need to be performed to the network, which is outside the scope the current work.

The rest of the paper is organized as follows. Section 2 briefly summaries related works, Section 3 provides an overview of our methodology, while Section 4 presents the results of our experiments. Finally, concluding remarks and future works are summarized in Section 5.

2 Related work

There is a clear tendency in the literature towards automatizing the design of NN. In AdaNet [1] the authors aim to optimize the network architecture and the internal weights, by balancing the trade-off between model complexity and empirical risk minimization. Although the claim is that networks discovered through this approach perform better than those found with a grid search, only limited results for binary classifiers trained on subsets of CIFAR-10 [4] dataset are presented.

Auto-Net [5] is a framework that automatically tunes feed-forward neural networks without human intervention. The authors focus only on fully-connected neural networks in order to keep the number of hyper-parameters at a manageable level. They tune 63 network and layer dependant hyper-parameters for networks which are at most 6 layers deep. With this method they generated an ensemble of 39 models that outperformed all human experts and was the first automatically generated model to win an image competition ³.

On the same track, a framework for large scale image classifiers [6] makes use of genetic intuitive mutations to explore unprecedented large search spaces. The fully automatic evolution culminates with a trained network that reaches 94.6% accuracy on CIFAR-10, 2% lower than manually engineered state-of-the-art networks. However, the prohibitive computational cost of the framework, that required $4 \cdot 10^{20}$ FLOP for a 12 layer network, makes this approach unfeasible in practical scenarios with bigger datasets.

Other works [7–9] that discover good combination of hyper-parameters do not employ neuro-evolution, therefore the discovery requirements in terms of time and resources are less expensive, but have other drawbacks such as limited search space, considerable loss in accuracy, limited usability, and need of retraining from scratch for a considerable amount of time after each network alteration.

³ <http://automl.chalearn.org/>

Other steps towards incremental training are presented in [10, 11], where the goal is to transfer knowledge from a small network towards a significantly larger network under some architectural constraints. These approaches have a twofold benefit. First, they explore the design space of current state-of-the-art networks in order to find better performers. Second, they avoid information loss when the network structure has to be slightly modified. Although they improve state-of-the-art results for the ImageNet classification task, it would be interesting to compare in terms of time/accuracy the networks obtained through transferred learning with the same networks trained from scratch.

3 Methodology

Let us consider a generic CNN \mathcal{N} composed of n layers. The incremental training method proposed in this work begins by partitioning the original network \mathcal{N} into K sub-networks \mathcal{S}_k , with $k = 1, \dots, K$ and $K \leq n$. In our current implementation, the partitioning is done a priori and does not change throughout the training process, although this will be subject of future studies. A sub-network can be composed of one or several layers, with the only constraint that at least one of the layers must contain trainable parameters; in other words, pooling and dropout cannot constitute a sub-network by themselves. The classifier block at the end of the network, which might include fully-connected layers or global pooling layers, is not considered in the K partitions.

Following the scheme illustrated in Fig. 1, the training process starts with sub-network \mathcal{S}_1 attached to the classifier block. To determine when is the optimal time to add the second sub-network \mathcal{S}_2 between \mathcal{S}_1 and the classifier, we compute every *window size* (ws) epochs the improvement in the validation accuracy. When the improvement observed is below a threshold we stop the training and we increase the network depth by adding the next sub-network. This process, which is illustrated in Fig. 2, repeats until all K sub-networks are incorporated or until another custom criteria is met.

Each time a new sub-network \mathcal{S}_{k+1} is inserted in the current architecture, its weights need to be initialized. This step is delicate, as our experiments show empirically that a non-optimal initialization (e.g., random) might prevent the network from reaching the same accuracy as the original one, or might slow down significantly the incremental training process (Section 4). To overcome this issue, we propose a more efficient initialization technique, that we refer to as *look-ahead*. Given a current network constituted by k sub-networks, our look-ahead method consists of training \mathcal{S}_{k+1} for a few epochs based on the input generated by the already trained uppermost part of the network. This strategy provides a more informed starting point for the new sub-network, since it looks ahead towards more complex features that can be obtained based on the previously learned ones. We remark that the weights of the original sub-network $\mathcal{S}_1 + \mathcal{S}_2 + \dots + \mathcal{S}_k$ are frozen during the look-ahead process that initializes \mathcal{S}_{k+1} ; this implies that no time is spent in back-propagating in that region of the network. Moreover, the depth of the look-ahead tends to be comparably smaller than the depth

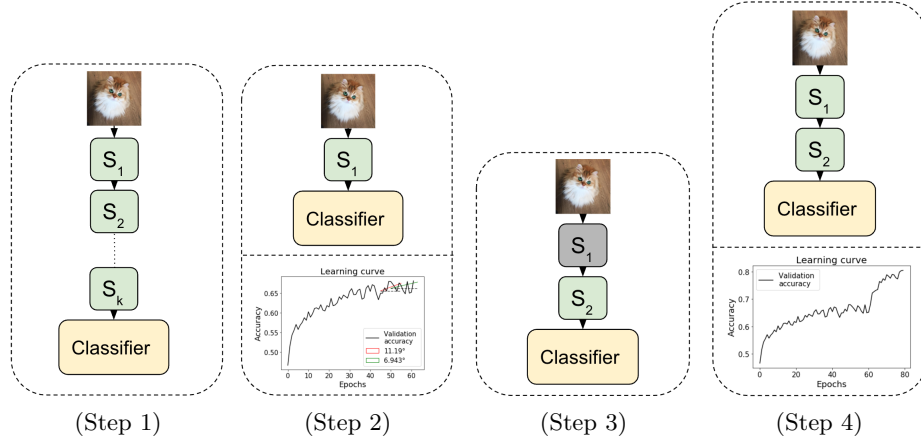


Fig. 1. A schematic description of incremental training for a CNN partitioned into K sub-networks. Step 1 depicts the original network that is already split into K sub-networks followed by the classifier layer. The incremental training process starts with the sub-network S_1 as shown in Step 2. The learning curve is monitored and when the termination criterion described in Section 3 is met, the training is stopped. S_2 is inserted in the current network in Step 3 and is trained only for a few epochs based on the frozen weights of S_1 . The obtained network $S_1 + S_2$ in Step 4 follows the same process as Step 2 and Step 3. The process finishes when either all the sub-networks are incorporated or when a custom criteria is met.




Fig. 2. Visualization of the termination criteria. Every *window size* (ws) epochs we compute the angle between the linear approximation of the last ws accuracy points and the x -axis. This angle shows the improvement of the validation accuracy. The training is stopped when $\alpha_i \leq \gamma \alpha_{i-1}$, where γ is a predefined threshold and α_i is the angle characterizing the accuracy for the i -th window.

of the final network, therefore the training of the look-ahead is not considered expensive.

Table 1. Configuration details adopted during regular and incremental training.

Panel A: Data pre-processing.		Panel B: Hyper-parameters.	
Parameter	Value	Parameter	Value
Train/Validation split	10%	Optimizer	RMSProp [14]
Width shift	4 px	Momentum	0.9
Height shift	4 px	Learning rate	10^{-4} (no decay)
Horizontal flip	Yes	Batch size	128 (ResNet-X)
			32 (VGG-16)
		Weight decay	10^{-4}
		Weight initialization	He [15]

4 Experiments

In this section, we train several state-of-the-art CNNs to compare the performance of our incremental training method, with the one of the classical algorithm. All runs involve single-precision arithmetic and are performed on IBM⁴ POWER8 Minsky compute nodes, equipped with four Nvidia P100 GPUs.

In all experiments we use the CIFAR-10 dataset. CIFAR-10 is a collection of 60,000 RGB images of 32x32x3 pixels each, evenly distributed across 10 mutually exclusive classes. On top of it, we also use the simple data augmentation described in [13] and summarized in Table 1, Panel A. Concerning the baseline networks, we chose ResNet [17] and VGGNet [16]: the former is tailored for CIFAR-10, while the latter is sized for more complex datasets.

The aim of our experiments is to demonstrate the advantages of the incremental training method compared to the regular one. Therefore, in our runs we do not employ all the advanced tricks (e.g., learning rate schedule) needed to reach the highest state-of-the-art accuracy. The values of the most relevant hyper-parameters are summarized in Table 1, Panel B. The partitioning of the networks for the incremental training is mainly based on simplicity. We group layers with the same number of filters (in the case of convolutions) and the same number of output nodes (in the case of dense layers). We tested different types of network partitioning observing no major impact in the overall performance.

In Fig. 3 we show a comparison between the regular and incremental training methods: for the latter, we plot results using both random and look-ahead initialization. The validation accuracy (first row) shows that the look-ahead outperforms the random initialization of the sub-networks. This behavior is better illustrated in the case of ResNet-56 network, in contrast with the behavior of VGGNet, as the over-parametrized network mitigates the outperformance. On the second row we estimate the cost per inference of each sub-network in terms of FLOPs. The plots highlight the benefit of our incremental training, that performs a much lower amount of calculations for the majority of the training time.

⁴ IBM, the IBM logo, ibm.com, OpenPOWER are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. Other product and service names might be trademarks of IBM or other companies.

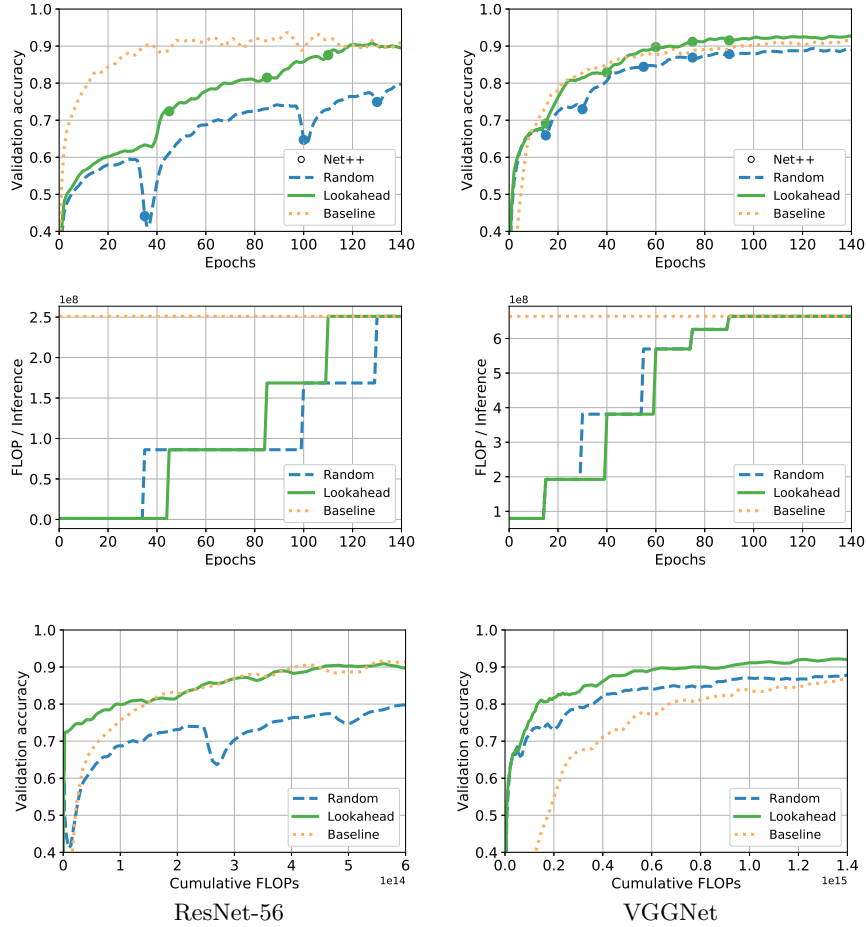


Fig. 3. Comparison in terms of accuracy (first row), inference cost (second row), and cumulative FLOPs (third row) between the training of the original network from scratch and its incremental counterpart on the CIFAR-10 dataset.

On the third and last row we observe that, while the regular version of the VGGNet reached 90% accuracy using 1.4 PFLOP, the incremental version needed 0.8 PFLOP and only 40% of the total number of parameters. In both networks the look-ahead training converged faster.

On top of these advantages, the incremental training allows to trace the importance of each sub-network during the training. For instance, in the case of VGGNet, the last two sub-networks bring a marginal improvement (less than 1% accuracy), while increasing by 16% the inference cost and by 80% the number of parameters. In Fig. 4 we show that only 5% of the VGGNet parameters are enough to reach 85% accuracy, while with 42% we arrive to the baseline accuracy. Therefore, the incremental training can be of great benefit for applications in

Table 2. Classification error on CIFAR-10 test set.

Network	Test error(%)		
	Baseline	Look-ahead	Random
VGGNet	90.06	90.50	88.91
ResNet-20	87.81	87.52	79.20
ResNet-32	88.93	88.90	79.89
ResNet-44	89.40	89.32	80.20
ResNet-56	89.51	89.56	80.96

which the last 1-2% of accuracy are not crucial and a slightly lower performance is within accepted limits. Indeed, it can dynamically determine to stop the network expansion early, saving a lot of computational resources without compromising accuracy.




Fig. 4. Validation accuracy as a function of the number of parameters in the network. The final accuracy obtained by the VGGNet (which consists of 35M parameters) is reached with only 15M parameters with the incremental training.




Fig. 5. Training of VGGNet on CIFAR-10 images resized to 224x224x3 pixels. For each sub-network we used the maximum possible batch size that fitted the GPU. This allowed for a larger learning rate for the first sub-networks, that led to a faster convergence than the baseline.

5 Conclusion

We proposed an incremental training method for CNNs. We demonstrated that our method reaches same or slightly better accuracy than regular training methods on state-of-the-art networks. This was achieved by using a more informed initialization of the network extension, which we call the look-ahead. Clear benefits of the incremental training are the faster convergence, the intuitive understanding of the importance of each sub-network in the overall performance, and the smooth synergy between training and optimal network depth discovery. In future works, we plan to generalize the present approach to other types of networks, such as Recurrent Neural Networks and Long Short Term Memory.

References

1. Corinna Cortes, Xavi Gonzalvo, Vitaly Kuznetsov, Mehryar Mohri, and Scott Yang. Adanet: Adaptive structural learning of artificial neural networks. *arXiv preprint arXiv:1607.01097*, 2016.
2. Jae-Yoon Jung and James A. Reggia. The automated design of artificial neural networks using evolutionary computation. In *Success in Evolutionary Computation*, pages 19–41. Springer, 2008.
3. Dong Yu, Li Deng, Frank Torsten Bernd Seide, Gang Li. Discriminative pretraining of deep neural networks, May 30 2013. US Patent App. 13/304,643.
4. Alex Krizhevsky and Geoffrey Hinton. Learning multiple layers of features from tiny images. 2009.
5. Hector Mendoza, Aaron Klein, Matthias Feurer, Jost Tobias Springenberg, and Frank Hutter. Towards automatically-tuned neural networks. In *Workshop on Automatic Machine Learning*, pages 58–65, 2016.
6. Esteban Real, Sherry Moore, Andrew Selle, Saurabh Saxena, Yutaka Leon Suenmatsu, Quoc Le, and Alex Kurakin. Large-scale evolution of image classifiers. *arXiv preprint arXiv:1703.01041*, 2017.
7. Junqi Jin, Ziang Yan, Kun Fu, Nan Jiang, and Changshui Zhang. Neural network architecture optimization through submodularity and supermodularity. *arXiv preprint arXiv:1609.00074*, 2016.
8. Song Han, Huizi Mao, and William J Dally. Deep compression: Compressing deep neural networks with pruning, trained quantization and Huffman coding. *arXiv preprint arXiv:1510.00149*, 2015.
9. Tobias Domhan, Jost Tobias Springenberg, and Frank Hutter. Speeding up automatic hyperparameter optimization of deep neural networks by extrapolation of learning curves. In *IJCAI*, pages 3460–3468, 2015.
10. Tianqi Chen, Ian Goodfellow, and Jonathon Shlens. Net2net: Accelerating learning via knowledge transfer. *arXiv preprint arXiv:1511.05641*, 2015.
11. Tao Wei, Changhu Wang, Yong Rui, and Chang Wen Chen. Network morphism. In *International Conference on Machine Learning*, pages 564–572, 2016.
12. Matthew D. Zeiler and Rob Fergus. Visualizing and understanding convolutional networks. In *European conference on computer vision*, pages 818–833. Springer, 2014.
13. Chen-Yu Lee, Saining Xie, Patrick Gallagher, Zhengyou Zhang, and Zhuowen Tu. Deeply-supervised nets. In *Artificial Intelligence and Statistics*, pages 562–570, 2015.
14. Tijmen Tieleman and Geoffrey Hinton. Lecture 6.5-rmsprop: Divide the gradient by a running average of its recent magnitude. *COURSERA: Neural networks for machine learning*, 4(2), 2012.
15. Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE international conference on computer vision*, pages 1026–1034, 2015.
16. Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
17. Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
18. Barret Zoph and Quoc V Le. Neural architecture search with reinforcement learning. *arXiv preprint arXiv:1611.01578*, 2016.