

Deep Cascade Learning

Enrique S. Marquez[✉], Jonathon S. Hare, and Mahesan Niranjan

Abstract—In this paper, we propose a novel approach for efficient training of deep neural networks in a bottom-up fashion using a layered structure. Our algorithm, which we refer to as deep cascade learning, is motivated by the cascade correlation approach of Fahlman and Lebiere, who introduced it in the context of perceptrons. We demonstrate our algorithm on networks of convolutional layers, though its applicability is more general. Such training of deep networks in a cascade directly circumvents the well-known vanishing gradient problem by ensuring that the output is always adjacent to the layer being trained. We present empirical evaluations comparing our deep cascade training with standard end–end training using back propagation of two convolutional neural network architectures on benchmark image classification tasks (CIFAR-10 and CIFAR-100). We then investigate the features learned by the approach and find that better, domain-specific, representations are learned in early layers when compared to what is learned in end–end training. This is partially attributable to the vanishing gradient problem that inhibits early layer filters to change significantly from their initial settings. While both networks perform similarly overall, recognition accuracy increases progressively with each added layer, with discriminative features learned in every stage of the network, whereas in end–end training, no such systematic feature representation was observed. We also show that such cascade training has significant computational and memory advantages over end–end training, and can be used as a pretraining algorithm to obtain a better performance.

Index Terms—Adaptive learning, cascade correlation, convolutional neural networks (CNNs), deep learning, image classification.

I. INTRODUCTION

DEEP convolutional networks have recently shown impressive results in a range of hard problems in AI, such as computer vision. However, there is still not clear understanding regarding how, and what, they learn. These models are typically trained end–end to capture low- and high-level features on every convolutional layer. There are still a number of problems with these networks that have yet to be overcome in order to obtain even better performance in computer vision tasks. In particular, one current community-wide trend is to build deeper and deeper networks; during training, these networks fall foul of an issue known as the vanishing gradient problem. The vanishing gradient problem manifests itself in

these networks because the gradient-based weight updates derived through the chain rule for differentiation are the products of n small numbers, where n is the number of layers being backward propagated through. In this paper, we aim to directly tackle the vanishing gradient problem by proposing a training algorithm that trains the network from the bottom-to-top layer incrementally, and ensures that the layers being trained are always *close* to the output layer. This algorithm has advantages in terms of complexity by reducing training time and can potentially also use less memory. The algorithm also has prospective use in building architectures without static depth that adapt their complexity to the data.

Several attempts have been proposed to circumvent complexity in learning. Platt [2] developed the Resource Allocating Network that allocates the memory based on the number of captured patterns, and learns these representations quickly. This network was then further enhanced by changing the LMS algorithm to include the extended Kalman filter, and by pruning and replacing it improved both in terms of memory and performance [3], [4]. Further, Shadafan *et al.* [5] present a sequential construction of multilayer perceptron (MLP) classifiers trained locally by recursive least squares algorithm. Compressing, pruning, and binarization of the weights in a deep model have also been developed to diminish the learning complexity of convolutional neural networks (CNNs) [6], [7].

In the late 1980s, Fahlman and Lebiere [1] proposed the cascade correlation algorithm/architecture as an approach to sequentially train perceptrons and connect their outputs to perform a single classification. Inspired by this idea, we have developed an approach to cascaded layerwise learning that can be applied to modern deep neural network architectures that we term deep cascade learning. Our algorithm reduces the memory and time requirements of the training compared with the traditional end–end backpropagation, and circumvents the vanishing gradient problem by learning feature representations that have increased correlation with the output on every layer.

Many of the core ideas behind CNNs occurred in the late 1970s with the neocognitron model [8], but failed to fully catch on for computational reasons. It was not until the development of LeNet-5 that CNNs took shape [9]. A great contribution to convolutional networks and an upgrade on LeNet style architectures came from generalizing the deep belief network idea [10] to a convolutional network [11]. However, with recent community-wide shift toward the use of very large models [12] (e.g., 19.4 M parameters) trained on very large data sets [13] (e.g., ImageNet with 1.4 M images) using extensive computational resources, we see a revolution in achievable performances as well as our thinking about such inference problems. The breakthrough of deep CNNs arrived

Manuscript received August 14, 2017; revised December 7, 2017; accepted January 30, 2018. Date of publication March 6, 2018; date of current version October 16, 2018. This work was supported by the Department of Electronics and Computer Science, University of Southampton. (*Corresponding author: Enrique S. Marquez.*)

The authors are with the Department of Electronics and Computer Science, University of Southampton, Southampton, SO17 1BJ U.K. (e-mail: esm1g14@soton.ac.uk).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TNNLS.2018.2805098

with the ImageNet competition winner 2012 AlexNet [14]. Since then, deep learning has constantly been pushing the state-of-the-art accuracy in image classification. The community has now been using these extensive convolutional networks architecture not only on classification problems, but in other computer vision and signal processing settings, such as object localization [15], semantic segmentation [16], face recognition and identification [17], speech recognition [18], and text detection [19]. Convolutional networks are very flexible because they are often trained as feature extractors and not only as classification devices. Furthermore, these nets can not only learn robust feature, but can also learn discriminative binary hash codes [20]. In any case, deep learning still has a long way to go in order to substantially outperform human level knowledge [13], [21].

Recently, networks have increased depth in order to capture low- and high-level features at different stages of the networks. A few years back, the deepest network was AlexNet with five convolutional layers and two dense layers, but now techniques such as the stochastic depth procedure [22] have used more than 1200 layers to increase the performance of the network. The rationale for these deeper networks is that more layers should capture better high-level features. However, when performing backpropagation on deep networks, because of the multiplicative effect of the chain rule, the magnitude of the gradient is greater on layers that are closer to the output, making the weight updates of the initial layers significantly smaller (layers that are closer to the input then learn at a slower rate). This issue is called the vanishing gradient problem, and it affects every network that is trained with any kind of backpropagation algorithm that has multiple weight layers.

Multiple algorithms have been proposed to overcome the vanishing gradient problem. Residual networks [12] are non-feedforward networks made of residual blocks, which are composed of convolutional layers, batch normalization [23], and a bypass connection that helps to alleviate the vanishing gradient problem. However, ResNets are equivalent to ensembles of shallow networks and do not fully overcome the vanishing gradient [24]. More recently, deep stochastic depth networks [22] combine the residual networks architecture with an extended version of dropout to again further solve the vanishing gradient problem, obtaining improvements of $\sim 1\%$ over ResNets.

The remainder of this paper is organized as follows. Section I-A explains the cascade learning algorithm and analyzes its advantages. Section I-B shows the results and discussion of two experiments performed on two architectures. Finally, Section II summarizes the findings, contributions, and potential further work of this paper.

A. Deep Cascade Learning Algorithm

In this section, we describe the proposed deep cascade learning algorithm and discuss the computational advantages of training in a layerwise manner. All the codes used to generate the results in this paper can be found in the GitHub repository available at <http://github.com/EnriqueSMarquez/CascadeLearning>.

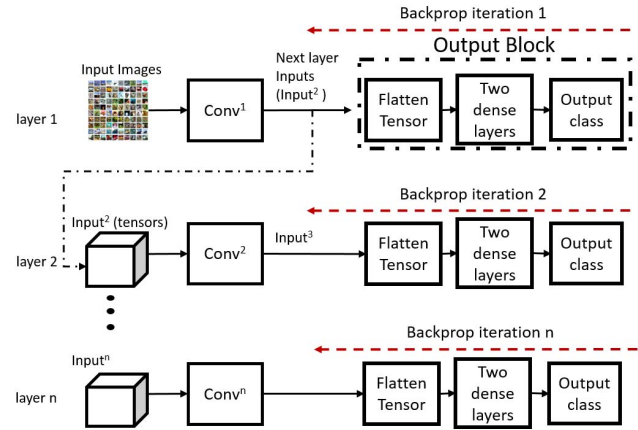


Fig. 1. Overview of deep cascade learning on a convolutional network with n layers. Input^i is the tensor generated by propagating the images through the layers up to and including Conv^{i-1} . Training proceeds layer by layer; at each stage using convolutional layer outputs as inputs to train the next layer. The features are flattened before feeding them to the classification stage. In contrast with the cascade correlation algorithm, the output block is discarded at the end of the iteration (see Algorithm 1), and typically, it contains a set of fully connected layers with nonlinearities and dropout.

1) *Algorithm Description:* As opposed to the cascade correlation algorithm, which sequentially trains perceptrons, we cascade layers of units. The proposed algorithm allows us to train deep networks in a cascade-like, or bottom-up layer-by-layer, manner. For the purposes of this paper, we focus on CNN architectures. The deep cascade learning algorithm splits the network into its layers and trains each layer one by one until all the layers in the input architecture have been trained, however, if no architecture is given, one can use the cascade learning to train as many layers as desired (e.g., until the validation error stabilizes). This training procedure allows us to counter the vanishing gradient problem by forcing the network to learn features correlated with the output on each and every layer. The training procedure can be generalized as “several” single-layer CNNs (subnetworks) that interconnect and can be trained one at a time from the bottom to top (see Fig. 1).

The algorithm takes as inputs the hyperparameters of the training algorithm (e.g., optimizer, loss, and epochs) and the model to train. Pseudocode of the cascade learning procedure can be found in Algorithm 1, and will be referred in further explanations of the algorithm. Learning starts by taking the first layer of the model and connecting it to the output with an “output block” (line 9), which might be several dense layers connected to the output [14], [26], or as it is sometimes shown in the literature, an average pooling layer and the output layer with an activation function [27]. Training using standard backpropagation then commences (using the presupplied parameters, loop in line 11) to learn weights for the first model layer (and the weights for the output block). Once the weights of the first layer have converged, the second layer can be learned by taking the second layer from the input model, connecting it to an output block (with the same form as for the first layer, but potentially a different dimensionality), and training it against the outputs with pseudoinputs created

Algorithm 1 Pseudocode of Cascade Learning Adapted From the Cascade Correlation Algorithm [25]. Training Is Performed in Batches, Hence Every Epoch Is Performed by Doing Backpropagation Through All the Batches of the Data

```

procedure CASCADE LEARNING(layers,  $\eta$ , epochs, epochsUpdate, out)
2:  Inputs  layers : model layers parameters (loss function, activation, regularization, number of filters, size of filters, stride)
            $\eta$  : Learning rate
           epochs : starting number of epochs
           k : epochs update constant
           out : output block specifications
3:  Output W : L layers with  $\mathbf{w}_l$  trained weights per layer
4:  for layer_index = 1 : L do                                     ▷ Cascading through trainable layers
   Init new layer and connect output block
5:  il  $\leftarrow$  epochs + k  $\times$  layer_index
6:  for i = 0; i++; i < il do                                     ▷ Loop through data il times
7:     $\mathbf{w}^{\text{new}} \leftarrow \mathbf{w}^{\text{old}} - \eta \nabla J(\mathbf{w})$                                ▷ Update weights by gradient descent
8:    if Validation error plateaus then
9:       $\eta \leftarrow \eta/10$                                                ▷ Change learning rate if update criteria is satisfied
10:   end if
11: end for
12: Disconnect output block and get new inputs
13: end for
end procedure

```

by forward propagating the actual inputs through the (fixed) first layer. This process can then repeat until all layers have been learned. At each stage, the pseudoinputs are generated by forward propagating the actual inputs through all the previously trained layers. It should be noted that once layer weights have been learned that they are fixed for all subsequent layers. Fig. 1 gives a graphical overview of the entire process.

Most hyperparameters in the algorithm remain the same across each layer, however, we have found it beneficial to dynamically increase the number of learning epochs as we get deeper into the network. Additionally, we start training the initial layers with orders of magnitude fewer epochs than we would if training end to end. The rationale for this is that each subnetwork fits the data faster than the end–end model and we do not want to overfit the data, especially in the lower layers. Overfitting in the lower layers would severely hamper the generalization ability of later layers. In our experiments, we have found that the number of epochs required to fit the data is dependable on the layer index, if a layer requires $i_{(\text{epochs})}$, the subsequent layer should require $i_{(\text{epochs})} + k$, where k is a constant whose value is set dependent on the data set.

A particular advantage of such cascaded training is that the backward propagated gradient is not diminished by hidden layers as happens in the end–end training. This is because every trainable layer is immediately adjacent to the output block. In essence, this should help the network obtain more robust representations at every layer. In Section I-B, we demonstrate this by comparing confusion matrices at different layers of networks trained using deep cascade learning and standard end–end backpropagation. The other advantages, as demonstrated in Section I-B, are that the complexity of learning is

reduced over end–end learning, both in terms of training time and memory.

2) *Cascade Learning as Supervised Pretraining Algorithm:* A particular appeal of deep neural networks is pretraining the weights to obtain a better initialization, and further achieve better minima. Starting from the work of Hinton *et al.* [10] on deep belief networks, unsupervised learning has been considered in the past as effective pretraining, initializing the weights which are then improved in a supervised learning setting. Although this was a great motivation, recent architectures [12], [21], [28], however, have ignored this and focused on pure supervised learning with random initialization.

The cascade learning can be used to initialize the filters in a CNN and diminish the impact of the vanishing gradient problem. After the weights have been pretrained using cascade learning, the network is tuned using traditional end–end training (both stages are supervised). When applying this procedure, it is imperative to reinitialize the output block after pretraining the network, otherwise the network would rapidly reach the suboptimal minimum obtained by the cascade learning. This does not provide better performance in terms of accuracy. In Section I-B3, we discuss how this technique may lead the network to better generalization.

3) *Time Complexity:* In a CNN, the time complexity of the convolutional layers is

$$O\left(\sum_{l=1}^d n_{l-1} s_l^2 n_l m_l^2 i\right) \quad (1)$$

where i is the number of training iterations, l is the layer index, d is the last layer index, n is the number of filters,

s and m is the size of the input and output (spatial size), respectively¹ [29].

Training a CNN using the deep cascade learning algorithm changes the time complexity as follows:

$$O\left(\sum_{l=1}^d n_{l-1} s_l^2 n_l m_l^2 i_l\right) \quad (2)$$

where i_l represents the number of training iterations for the l th layer. The main difference between both equations is the number of epochs for every layer, in (1) i is constant, while in (2) i depends on the layer index. Note in this analysis, we have purposefully ignored the cost of performing the forward passes to compute the pseudoinputs as this is essentially “free” if the algorithm is implemented in two threads (given in the following). The number of iterations in the cascade algorithm depends on the data set and the model architecture. The algorithm proportionally increases the number of epochs on every iteration since the early layers must not be overfit, while the later layers should be trained to more closely fit the data. In practice, as shown in the simulations (Section I-B), one can choose each i_l such that $i_1 \ll i$ and $i_L \leq i$, and obtain almost equivalent performance to the end–end trained network in a much shorter period of time. If $\sum_{l=1}^d i_l = i$, the time complexity of both training algorithms is the same, noting that improvements coming from caching the pseudoinputs are not considered.

There are two main ways of implementing the algorithm. The best and most efficient approach is by saving the pseudoinputs on disk once they have been computed; in order to compute the pseudoinputs for the next layer, only one has to forward propagate the cached pseudoinputs through a single layer. An alternate, naive, approach would be implementing the algorithm using two threads (or two GPUs), with one thread using the already trained layers to generate the pseudoinputs on demand and the other thread training the current layer. The disadvantage of this is that it would require the input to be forward propagated on each iteration. The first approach can further drastically decrease the runtime of the algorithm and the memory required to train the model at the expense of disk space used for storing cached pseudoinputs.

4) *Space Complexity*: When considering the space complexity and memory usage of a network, we not only consider both the number of parameters of the model, but also the amount of data that needs to be in memory in order to perform training of those parameters. In standard end–end backpropagation, intermediary results (e.g., response maps from convolutional layers and vectors from dense layers) need to be stored for an iteration of backpropagation. With modern hardware and optimizers (based on variants of minibatch stochastic gradient descent), we usually consider batches of data being used for the training, so the amount of intermediary data at each layer is multiplied by the batch size.

Aside from offline storage for caching pseudoinputs and storing trained weights, the cascade algorithm only requires that the weights of a single model layer, the output block

¹Note that this is the time complexity of a single forward pass; training increases this by a constant factor of about 3.

TABLE I

SPACE COMPLEXITY OF END–END TRAINING OF VARIOUS DEPTHS OF VGG STYLE NETWORKS. THE NUMBER OF PARAMETERS INCREASES WITH DEPTH. THE DATA STORAGE UNITS OF THE TRAINING DEPEND ON THE COMPUTATIONAL PRECISION

model	parameters	data storage	total
VGG-16	13.9M	311178	14.2M
VGG-19	19.1M	337802	19.5M
VGG-22	24.5M	364426	24.8M
VGG-25	29.8M	391050	30.2M
VGG-28	35.1M	417674	35.5M

weights, and the pseudoinputs of the current training batch are stored in RAM (on the CPU or GPU) at any one time. This *potentially* allows memory to be used much more effectively and allows models to be trained whose weights exceed the amount of available memory, however, this is drastically affected by the choice of output block architecture, and also the depth and overall architecture of the network in question.

To explore this further, consider the parameter and data complexity of a Visual Geometry Group Net (VGG-style network) of different depths. Assume that we can grow the depth in the same way as going between the VGG-16 and VGG-19 models in the original paper [26] (note, we are considering Model D in the original paper to be VGG-16 and Model E to be VGG-19), whereby to generate the next deeper architecture, we add an additional convolutional layer to the last three blocks of similarly sized convolutions. This process allows us to define models VGG-22, VGG-25, VGG-28, etc. The number of parameters and training memory complexity of these models is shown in Table I. The numbers in this table were computed on the assumption of a batch size of 1, input size of 32×32 , and the output block (last three fully connected/dense layers) consisting of 512, 256, and 10 units, respectively. The remainder of the model matches the description in the original paper [26], with blocks of 64, 128, 256, and 512 convolutional filters with a spatial size of 3×3 and the relevant max pooling between blocks. For simplicity, we assume that the convolutional filters are zero-padded so the size of the input does not diminish.

The key point to note from Table I is that as the model gets bigger, the amount of memory required for both parameters and for data storage of end–end training increases. With our proposed cascade learning approach, this is not the case; the total memory complexity is purely a function of the most complex cascaded subnetwork (network trained in one iteration of the cascade learning). In the case of all the above VGG-style networks, this happens very early in the cascading process. More specifically, this happens when cascading the second layer, as can be seen in Table II, which illustrates that after the second layer (or more concretely after the first max pooling), the complexity of subsequent iterations of cascading reduces. The assumption in computing the numbers in Table II is that the output blocks mirrored those of the end–end training had 512, 256, and 10 units, respectively.

If we consider Tables I and II together, we can see with the architectures in question that for smaller networks the end–end training will use less memory (although it is slower), while

TABLE II

SPACE COMPLEXITY OF CASCADE TRAINING OF VARIOUS LAYERS OF A VGG STYLE NETWORK. THE NUMBER OF PARAMETERS DECREASES WITH DEPTH. THE DATA STORAGE UNITS OF THE TRAINING DEPEND ON THE COMPUTATIONAL PRECISION

trainable layer #	parameters	data storage	total
1	33.6M	69386	33.7M
2	8.6M	131850	8.6M
<i>Pooling</i>			
3	16.9M	49930	17.0M
4	4.5M	66314	4.6M
<i>Pooling</i>			
5	8.8M	25354	8.8M
6	2.8M	25354	2.9M
...			

for deeper networks, the cascading algorithm will require less peak memory while bringing time complexity reductions. Given that the bulk of the space complexity for cascading comes as a result of the potentially massive number of trainable parameters in connecting the feature maps from the early convolutional layers to the first layer of output block, an obvious question is could we change the output block specification to reduce the space complexity for these layers? While not the key focus of this paper, initial experiments described in Section I-B1a) start to explore the effect of reduced complexity output blocks on overall network classification performance.

B. Experiments

The first experiment was performed on a less complex backpropagation problem and not on a CNN as explained in Section I-A. We decided to execute this experiment to quickly determine the efficiency of cascade learning. In this case, we have chosen a small three hidden layers MLP applied on the flattened MNIST data set. The results show that this algorithm is feasible and can obtain better generalization in early stages of the network with small improvements ($\sim 0.5\%$). This was a preliminary experiment, details can be found in the GitHub repository.

To demonstrate the effectiveness of the deep cascade learning algorithm, we apply it to two widely known architectures: a “VGG-style” network [26] and the “All-CNN” [27]. We have chosen these architectures for several reasons. First, they are still extensively used in the computer vision community, and second, they inspired state-of-the-art architectures, such as ResNets and FractalNets. Explicitly, the VGG net shows how small filters (3×3) can capture patterns at different scales by just performing enough subsampling. The All-CNN gave the idea of performing the subsampling with an extra convolutional layer rather than a pooling layer, and performs the classification using global average pooling and a dense layer to diminish the complexity of the network. The representations learned in each layer through end–end training are compared to the ones generated by deep cascade learning. In order to make a layerwise comparison, we first train an end–end model, and then use the already trained filters to train classifiers by attaching and training output blocks (the model layer weights are fixed at this point, however, in contrast

to cascade learning). The training parameters of the models remain as similar as possible to make a fair comparison.

The learning rate in both experiments is diminished when the validation error plateaus. Taking into account that the data set is noisy and the error does not necessarily decrease after every epoch, we evaluate the performance after each epoch to determine whether the learning rate should be changed. More specifically, we use a mean-window approach that computes the average of the last five epochs and the last 10 epochs, and if the difference is negative then the learning rate is decreased by a factor of 10. The size of the window was tuned for the cascade learning only; if this approach is used in other training procedures, it might be necessary to increase the size of the window.

The increase in the epochs in the cascade algorithm varies depending on the data set. We performed experiments with an initial number of epochs ranging from 10 to 100 without any real change in the overall result, hence, 10 epochs as starting point is the most convenient. In all the experiments presented here, every layer iteration initializes a new output block, which in this case consists of two dense layers with ReLu activation [30]. The number of neurons in the first layer will depend on the dimensionality of the input vector, it may vary between 64 and 512 units, the second layer contains half as many units as in the first layer. The final layer uses softmax activation and 10 or 100 units depending on the data set.

Data Sets: We have performed experiments using the CIFAR-10 and CIFAR-100 [31] image classification data sets, which have 10 and 100 target labels, respectively. Both data sets contain 60000 RGB 32×32 images split in three sets: 45000 images for training, 5000 images for validation, and 10000 images for testing. In our experiments, the data sets were normalized and whitened, however, we performed no further data augmentation, similar to the stochastic depth procedure [22].

1) CIFAR-10:

VGG Style Networks: The VGG network uses a weight decay of 0.001, and stochastic gradient descent with a starting learning rate of 0.01. Our VGG model contains six convolutional layers, starting with $128 \ 3 \times 3$ filters and duplicating them after a MaxPooling layer. The initial weights remained the same in the networks trained by the two approaches to make the convergence comparable.

a) Space complexity and output block specifications:

In order to test the memory complexity of this network we must take into account the output block specifications. Specifically, we must consider the first fully connected layer, which in most networks contains the biggest number of trainable parameters, particularly when connected to an early convolutional layer (see Section I-A4). On the first iteration of cascade learning, the output is $128 \times 32 \times 32$, hence, the number of neurons (n) in the first fully connected layer must be small enough to avoid running out of memory, but without jeopardizing robustness in terms of predictive accuracy. We have performed an evaluation by cascading this architecture with output blocks with a range of different parameter complexities. Table III shows the number of parameters of every layer as well as the performance for output blocks with first fully connected




Fig. 2. Time complexity comparison between cascade learning, end-end, and pretrained using cascade learning (see Section I-B3 for details and results on pretraining using cascade learning). Multiple VGG networks were executed within a range of starting number of epochs (10–100) (left) and depth (3,6,9) (right). Black pentagons: runs executing the naive approach for both cascade learning and the pretraining stage. Blue solid dots: optimal run, which caches the pseudoinputs after every iteration.

TABLE III

PARAMETER COMPLEXITY COMPARISON USING DIFFERENT OUTPUT BLOCK SPECIFICATIONS SHOWS THE EFFECT OF USING BETWEEN 64 AND 512 UNITS IN THE FIRST FULLY CONNECTED LAYER (WHICH IS MOST CORRELATED WITH THE COMPLEXITY). LEFT: NUMBER OF PARAMETERS. RIGHT: ACCURACY. BOTTOM ROW SHOWS THE PARAMETERS COMPLEXITY OF THE END-END MODEL. THE INCREASE IN MEMORY COMPLEXITY ON EARLY STAGES CAN BE NAIVELY REDUCED BY DECREASING n . POTENTIALLY, MEMORY REDUCTION TECHNIQUES ON THE FIRST FULLY CONNECTED LAYER ARE APPLICABLE AT EARLY STAGES OF THE NETWORK. LATER LAYERS ARE LESS COMPLEX

Training Regime.	Iter.	First output block unit count							
		64		128		256		512	
		param.	acc.	param.	acc.	param.	acc.	param.	acc.
CL	1	8.4e6	0.63	1.7e7	0.64	3.4e7	0.66	6.7e7	0.66
	2	8.5e6	0.69	1.7e7	0.72	3.4e7	0.72	6.7e7	0.73
	3	2.5e6	0.77	4.4e6	0.78	8.6e6	0.79	1.7e7	0.80
	4	4.5e6	0.80	8.7e6	0.81	1.7e7	0.81	3.4e7	0.81
	5	4.8e6	0.82	9.0e6	0.83	1.7e7	0.83	3.4e7	0.83
	6	1.6e6	0.83	2.7e6	0.84	4.8e6	0.84	9.1e6	0.84
End-End		2.8e6	0.86	3.9e6	0.87	6.0e6	0.87	1.0e7	0.87

layer sizes of $n = \{64, 128, 256, 512\}$. In terms of parameters, cascade learning for early iterations can require more space than the entire end-end network unless the overall model is deep. The impact of this disadvantage can be overcome by choosing a smaller n , and as shown in Table III, the hit on accuracy need not be particularly high when compared to the reduction in parameters and saving of memory.

Reducing the number of units can efficiently diminish the parameters of the network. However, in cases where the input image is massive, more advanced algorithms to counter the exploding number of parameters are applicable, such as tensorizing neural networks and hashed nets [32], [33]. Based on their findings, applying those types of transformations to the first fully connected layer should not affect the results.

b) Training time complexity and relationship with depth and starting number of epochs: Equation 2 is dependent on the starting number of epochs i_l and its proportionality with depth. In Fig. 2, we explored the effects of the time complexity by these two variables. To reproduce Fig. 2 (left), several networks were cascaded with $i_l = [10, 30, 50, 70, 100]$, the overall required time is not drastically affected by i_l . For this particular experiment if $i_l > 50$, each iteration is more likely to be stopped early due to overfitting. Fig. 2 (right) shows the results on a similar experiment with varying network




Fig. 3. Comparison of confusion matrices in a VGG network trained using the cascade algorithm and the end-end training on CIFAR-10. First two layers of the end-end training do not show correlation with the output. While accuracy increases proportionally with the number of layer using the cascade learning, it shows more stable features at every stage of the network.

depth ($d = [3, 6, 9]$). Cascading shallow networks outperforms end-end training in terms of time. The epochs update constant (k in Section I-A1) should be minimized on deeper networks to avoid an excessive overall number of epochs. Fig. 2 shows the importance of caching the pseudoinputs, the black pentagons (naive run) are shifted to the right in relation to solid blue dots (enhanced run).

Fig. 3 shows confusion matrices from both algorithms across the classifiers trained on each layer. In this experiment, we found that the features learned using the cascade algorithm are less noisy, and more correlated with the output in most




Fig. 4. Comparison of confusion matrices in All-CNN network trained using the cascade algorithm and the end-end training on CIFAR-10. Features learned by cascading the layers are less noisy, and more stable.




Fig. 5. Visualization of the filters learned in the first layer in both algorithms. (a) Cascade learning. (b) End-end. Each patch corresponds to one 3×3 filter. Filters learned using the cascade learning show more clear representations with a wide number of rotations, while in the end-end most filters are redundant.

stages of the network. The results of the experiment shows that the features learned using the end-end training in the first and second layer are not correlated with the output; in this case, the trained output block classifier always makes the same prediction, and hence, the feature vector is not sparse at all. The third layer starts building the robustness of the features with an accuracy of 67.2%, and the peak is reached in the last layer with 85.6%. In contrast, with the cascade learning, discriminative features are learned on every layer of the network. At the third layer, classes such as airplane and ship are strongly correlated with the features generated in both cases. The end-end training mostly fails to generalize correctly in classes related to animals.

On every iteration of the cascade algorithm, the subnetworks have a tendency to overfit the data. However, this is not entirely a problem as we have found that overfitting mostly occurs in the dense layers connected in the output block, and those are not part of the resulting model. In this way, we avoid generating overfitted pseudoinputs for the next iteration, hence disconnecting the dense layers works as a matter of regularization.

One of the ways of determining if the vanishing gradient problem has been somehow diminished is by observing the filters/weights on the first layer (the most affected one by this issue). If the magnitude of the gradient in the first layer is small, then the filters do not change much from the initialized one. Fig. 5 shows the filters learned using both algorithms. The cascade algorithm learned a range of different filters with different orientation and frequency responses, while using an end-end training the filters learned are less representative. Some filters in the end-end training are overlapping, this generates a problem since the information that is being captured is redundant.

It is naive to assume the problem is alleviated because the filters on the cascade learning are further apart from the initial filters. Hence, to complement the visualization of the filters, we calculated the magnitude of the gradient after

every mini-batch forward pass on both cascade learning and end-end and plotted the results on Fig. 6. For the end-end training, the gradient was computed at every convolutional layer for all the epochs. For the cascade learning, the gradients were calculated on every iteration on the core correspondent convolutional layer. The curves are generated by averaging the mini-batch gradients in each epoch.

In contrast with cascade learning, the magnitude of the gradients of end-end training, on early layers, is significantly smaller than those on deeper layers. Overall, the gradients are higher for the cascade learning. Cascade learning requires fewer epochs with high updates on the weights to quickly fit the data on every iteration. With end-end training the opposite occurs; it requires more epochs (because of the small updates) to fit the data.

All-CNN: This architecture contains only convolutional layers, the downsampling is performed using a convolutional layer with stride of 2 rather than a pooling operation. It also performs the classification by downsampling the image until the dimensionality of the output matches the targets. The All-CNN paper [27] describes three model architectures. We have performed our experiments using model C that contains seven core convolutional layers, and four 1×1 convolutional layers to perform the classification with an average pooling and softmax layers as the output block. In this case where the output block contains an average pooling and a softmax activation, each layer would learn the filters required to classify the data and not to generate robust filters. Hence, to make a fair comparison of the filters we have changed the output block of the All-CNN to three dense layers with softmax activation at the end. In the All-CNN report, it is stated that changing the output block may results in a decrease of the performance, however, in this study we aim to fairly compare both algorithms in every stage rather than final classification result. The parameters used when cascading this architecture varies between 2.7×10^6 and 0.33×10^6 ; on the other hand the end-end training requires us to store 1.3×10^6 parameters.

The All-CNN, using an end-end training, learns better representations on early layers than the VGG-style net. The first convolutional layer achieves a performance in the orders of 20% by learning three classes at the most, this can be observed in the confusion matrix in Fig. 4. In contrast with the end-end training, the accuracy when cascading this architecture progressively increases with the iterations, learning discriminative representations at every stage of the network going from 65% to 83.4%.

Fig. 7 compares the performance of both algorithms on each layer. The accuracy in the cascade learning increases with the number of layers. In addition, the variance of the performance is very low in comparison with the end-end, because it forces the network to learn similar filters in every run, decreasing the impact of a poor initialization.

We have found that for a given iteration more than 50 epochs are not necessary to achieve a reasonable level of accuracy without overfitting the data. We also tested the time complexity of this model within a range of starting epochs (similar experiment in Section I-B1). We tested the time




Fig. 6. Magnitude of the gradient after every minibatch forward pass on the convolutional layers of the end-end training (right) and the concatenated gradients of every cascade learning (left) iteration. Vertical lines: start of a new iteration. Curves were smoothed by averaging the gradients (of every batch) on every epoch.




Fig. 7. Performance on every layer on both architectures. Top: VGG. Bottom: All-CNN. Cascade learning has a lower variance making the initialization less relevant to the classification at each layer. It also shows a progressive increase in the performance without the fluctuations presented in the end-end training.

complexity from 10 starting epochs to 50 (epochs increase by 10 on every iteration with a ceiling on 50). The time complexity for the All-CNN model C is reduced by ~ 2.5 regardless of the starting number of epochs.

2) *CIFAR-100*: Similar to the previous experiments, we have tested how the cascade algorithm behaves with a 100-class problem using the CIFAR-100 data set [31]. The experimental settings remain the same as Section I-B1, and the main change to the model is that the output layer now has 100 units to match the number of classes.

In a VGG-style network, the comparison between both algorithms is similar to a 10 class problem. In end-end




TABLE IV

COMPARISON OF ACCURACY PER LAYER USING THE CASCADE ALGORITHM AND END-END TRAINING ON CIFAR-100 IN BOTH ARCHITECTURES. USING THE CASCADE LEARNING OUTPERFORMS ALMOST ALL THE LAYERS IN A VGG NETWORK, AND ALMOST ACHIEVES THE SAME ACCURACY IN THE FINAL STAGE. THE ALL-CNN WITH AN END-END TRAINING OUTPERFORMS IN THE FINAL CLASSIFICATION, HOWEVER, THE FIRST THREE LAYERS DO NOT LEARN STRONG CORRELATIONS LIKE WHEN USING THE CASCADE LEARNING

		1	2	3	4	5	6	7	Pre
VGG	CL	0.35	0.39	0.50	0.50	0.53	0.59	-	0.63
	E-E	0.01	0.03	0.22	0.14	0.35	0.60	-	
The All CNN	CL	0.31	0.39	0.47	0.46	0.49	0.54	0.52	0.67
	E-E	0.03	0.05	0.03	0.41	0.54	0.61	0.62	

training, the first two layers do not learn meaningful representations, and each layer learns better features using the cascade algorithm. However, the end-end training performs better by 1% on the final classification.

In the All-CNN network, the features learned in the end-end model remained more stable than in CIFAR-10. Similarly, than in Section I-B1, the first four layers were outperformed by the cascaded model. The cascade network in overall had better performance in the end-end model by 6% and 10% on the last layers.

The results on a 100-class problem are arguably the same as in a 10-class one. It is noted that the All-CNN network, when trained end-end, can outperform the cascade algorithm in the final classification but not in the early layers. In the VGG-style network, deep cascade learning build more robust features in all the layers, except for the last layer which had a difference of 1%. Table IV shows a summary of the results on every layer for both algorithms.

3) *Pretraining With Cascade Learning*: In the experimental work described so far, the main advantages of cascade learning come from: 1) reduced computation, albeit at the loss of some performance in comparison to end-end training and 2) a better representation at intermediate layers. We next sought to explore if the representations learned by the computationally efficient cascading approach could form good

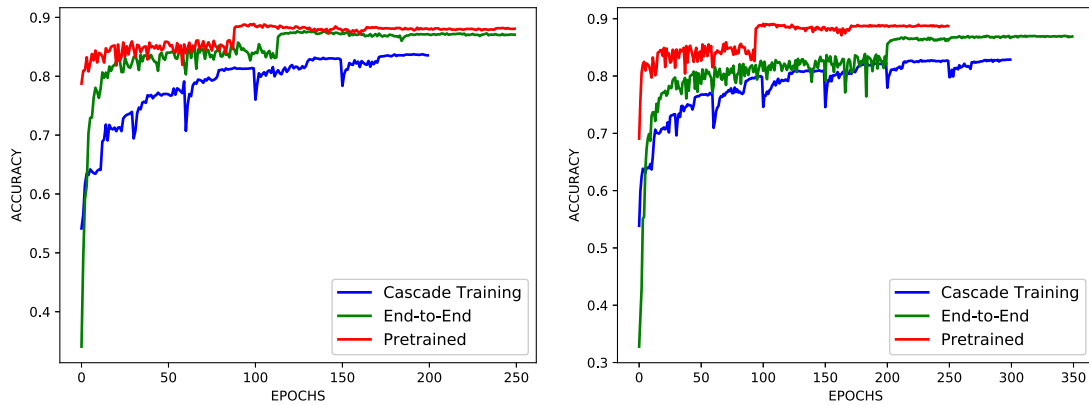


Fig. 8. Performance comparison on CIFAR-10 between pretrained network and random initialization. Left: VGG. Right: All-CNN. The step bumps in the cascade learning are generated due to the start of a new iteration or changes in the learning rate.




Fig. 9. Filters on the first layer for at different stages of the procedure on the VGG network defined in Section I-B1. (a) Initial random weights. (b) End-end. (c) Cascaded. (d) End-end trained network initialized by cascade learning.

initializations of end-end trained networks and achieve performance improvements.

The weights are initialized randomly. Then the procedure is divided into two stages, first, we cascade the network with the minimum number of epochs to diminish its time complexity. Finally, the network is fine-tuned using a backpropagation and stochastic gradient descent, similar to the end-end training. We applied this technique using a VGG-style network and the All-CNN. For more details on the architectures, refer to Section I-B1.

Fig. 8 shows the difference in performance given random and cascade learning initialization. The learning curves in Fig. 8 are for the VGG and the All-CNN architectures trained on CIFAR-10. The improvements in testing accuracy varies between $\sim 2\%$ and $\sim 3\%$ for the experiments developed in this section. However, the most interesting property comes as a consequence of the variation of the resulting weights after executing the cascade learning. As shown in Section I-B this variation is significantly smaller in contrast with its end-end

counterpart. Hence, the results obtained after preinitializing the network are more stable and less affected by poor initialization. Results on Fig. 2 show that even including the time of the tuning training stage, the time complexity can be reduced if the correct parameters for the cascade learning are chosen. It is important to mention that, the end-end training typically requires up to 250 epochs, while the tuning stage may only require a small fraction since the training is stopped when the training accuracy reaches ~ 0.999 .

The filters generated by the cascade learning filters are slightly overfitted (the first layer typically achieves $\sim 60\%$ on the unseen data and $\sim 95\%$ on the training data) as opposed to the end-end training, on which the filters are more likely to be close to its initialization. By pretraining with cascade learning, the network learns filters that are in between both scenarios (under and overfitness), this behavior can be spotted on Fig. 9.

Fig. 10 shows the test accuracy during training of a cascaded pretrained VGG model on CIFAR-100. Improvements

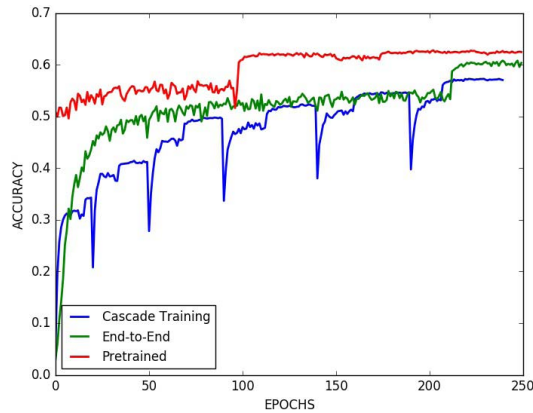


Fig. 10. Performance comparison between pretrained network and random initialization on CIFAR-100 using a VGG network.

of $\sim 2.5\%$ were achieved in the final classification. More details on this experiment are available in the GitHub repository accompanying this paper.

II. CONCLUSION

In this paper, we have proposed a new supervised learning algorithm to train deep neural networks. We validate our technique by studying an image classification problem on two widely used network architectures. The vanishing gradient problem is diminished by our deep cascade learning, because it is focused on learning more robust features and filters in early layers of the network. In addition, the time complexity is reduced because it no longer needs to forward propagate the data through the already trained layers on every epoch. In our experiments, the memory complexity is decreased more than three times for the VGG style network and four times for the All-CNN. Standard end-end training has a high variance in the performance, meaning that the initialization plays an important role in ensuring a good minimum is reached by each layer. Deep cascade learning generates a more stable output on every stage by learning similar representations at every run. In addition, the cascade learning algorithm has demonstrated to scale in 10 and 100 class problems, and shows improvements in the features that are learned across the stages of the network. Using this algorithm allows us to train deeper networks without the need to store the entire network in memory. It should be noted that our algorithm is not aimed at obtaining better performance than standard approaches, but with significant reduction in the memory and time requirements. We have shown that if improvements in generalization are expected, this algorithm has to be used as a pretraining algorithm technique.

There are many questions that are still yet to be answered. How deep can this algorithm go without losing robustness? We believe that if the performance cannot be improved by appending a new convolutional layer, l , it should at least be as good as in the previous layer, $l - 1$, by learning filters that directly map the input to the output (filters with 1 in the center, and zero in the borders). This might not happen because the layer might quickly find a local minimum. This could be avoided with a different type of initialization; most probably

one specialized for this algorithm. Our immediate next steps include observing how deep can the cascading algorithm can go without losing performance, similar to the experiment performed with deep residual network [12] and fractal networks [28], in order to measure to what extent the vanishing gradient problem is solved. In [12], the accuracy diminished when they went beyond 1200 layers, and hence the vanishing gradient problem was not entirely circumvented. We believe this algorithm might be able to go deeper without losing performance by partially overcoming the vanishing gradient problem, learning “mapping” filters to maintain the features sparseness, and learn a bigger set of high-level features. In addition, the deep cascade learning has the potential to find the number of layers required to fit a certain problem (adaptive architecture), similar to the cascade correlation [1], infinite restricted Boltzmann machine [34], and AdaNet [35].

ACKNOWLEDGMENT

The authors would like to thank the support of the NVIDIA Corporation with the donation of the Titan X GPU used for this research. They would also like to thank reviewers for their insightful comments and suggestions, as these led them to an improvement of this paper.

REFERENCES

- [1] S. E. Fahlman and C. Lebiere, “Advances in neural information processing systems,” in *The Cascade-Correlation Learning Architecture*, D. S. Touretzky, Ed. San Francisco, CA, USA: Morgan Kaufmann Publishers, 1990, pp. 524–532. [Online]. Available: <http://dl.acm.org/citation.cfm?id=109230.107380>
- [2] J. Platt, “A resource-allocating network for function interpolation,” *Neural Comput.*, vol. 3, no. 2, pp. 213–225, 1991.
- [3] V. Kadirkamanathan and M. Niranjan, “A function estimation approach to sequential learning with neural networks,” *Neural Comput.*, vol. 5, no. 6, pp. 954–975, Nov. 1993.
- [4] C. Molina and M. Niranjan, “Pruning with replacement on limited resource allocating networks by F-projections,” *Neural Comput.*, vol. 8, no. 4, pp. 855–868, May 1996.
- [5] R. Shadafan, *Sequential Training of Multilayer Perceptron Classifiers*. Cambridge, U.K.: Univ. Cambridge, 1995. [Online]. Available: <https://books.google.co.uk/books?id=HFtGwAACAAJ>
- [6] S. Han, H. Mao, and W. J. Dally, “Deep compression: Compressing deep neural network with pruning, trained quantization and Huffman coding,” *CoRR*, vol. abs/1510.00149, 2015. [Online]. Available: <http://arxiv.org/abs/1510.00149>
- [7] M. Rastegari, V. Ordonez, J. Redmon, and A. Farhadi. (2016). “XNOR-Net: ImageNet classification using binary convolutional neural networks.” [Online]. Available: <https://arxiv.org/abs/1603.05279>
- [8] K. Fukushima, “Neocognitron: A self-organizing neural network model for a mechanism of pattern recognition unaffected by shift in position,” *Biol. Cybern.*, vol. 36, no. 4, pp. 193–202, 1980. [Online]. Available: <http://dx.doi.org/10.1007/BF00344251>
- [9] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, “Gradient-based learning applied to document recognition,” *Proc. IEEE*, vol. 86, no. 11, pp. 2278–2324, Nov. 1998.
- [10] G. E. Hinton, S. Osindero, and Y.-W. Teh, “A fast learning algorithm for deep belief nets,” *Neural Comput.*, vol. 18, no. 7, pp. 1527–1554, 2006.
- [11] H. Lee, R. Grosse, R. Ranganath, and A. Y. Ng, “Convolutional deep belief networks for scalable unsupervised learning of hierarchical representations,” in *Proc. 26th Annu. Int. Conf. Mach. Learn.*, 2009, pp. 609–616.
- [12] K. He, X. Zhang, S. Ren, and J. Sun. (2015). “Deep residual learning for image recognition.” [Online]. Available: <https://arxiv.org/abs/1512.03385>
- [13] O. Russakovsky *et al.*, “ImageNet large scale visual recognition challenge,” *Int. J. Comput. Vis.*, vol. 115, no. 3, pp. 211–252, Dec. 2015.
- [14] A. Krizhevsky, I. Sutskever, and G. E. Hinton, “Imagenet classification with deep convolutional neural networks,” in *Proc. Adv. Neural Inf. Processing Syst.*, 2012, pp. 1097–1105.

- [15] R. Girshick, J. Donahue, T. Darrell, and J. Malik, "Rich feature hierarchies for accurate object detection and semantic segmentation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2014, pp. 1–8.
- [16] J. Long, E. Shelhamer, and T. Darrell, "Fully convolutional networks for semantic segmentation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2015, pp. 3431–3440.
- [17] Y. Taigman, M. Yang, M. Ranzato, and L. Wolf, "DeepFace: Closing the gap to human-level performance in face verification," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2014, pp. 1701–1708.
- [18] O. Abdel-Hamid, A.-R. Mohamed, H. Jiang, and G. Penn, "Applying convolutional neural networks concepts to hybrid NN-HMM model for speech recognition," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process. (ICASSP)*, Mar. 2012, pp. 4277–4280.
- [19] C. Yan, H. Xie, S. Liu, J. Yin, Y. Zhang, and Q. Dai, "Effective uyghur language text detection in complex background images for traffic prompt identification," *IEEE Trans. Intell. Transp. Syst.*, vol. 19, no. 1, pp. 220–229, Jan. 2018.
- [20] C. Yan, H. Xie, D. Yang, J. Yin, Y. Zhang, and Q. Dai, "Supervised hash coding with deep neural network for environment perception of intelligent vehicles," *IEEE Trans. Intell. Transp. Syst.*, vol. 19, no. 1, pp. 284–295, Jan. 2018.
- [21] K. He, X. Zhang, S. Ren, and J. Sun, "Delving deep into rectifiers: Surpassing human-level performance on imagenet classification," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Dec. 2015, pp. 1026–1034.
- [22] G. Huang, Y. Sun, Z. Liu, D. Sedra, and K. Q. Weinberger. (2016). "Deep networks with stochastic depth." [Online]. Available: <https://arxiv.org/abs/1603.09382>
- [23] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," in *Proc. 32nd Int. Conf. Mach. Learn.*, 2015, pp. 448–456.
- [24] A. Veit, M. J. Wilber, and S. Belongie, "Residual networks behave like ensembles of relatively shallow networks," in *Proc. Adv. Neural Inf. Process. Syst.*, 2016, pp. 550–558.
- [25] R. O. Duda, P. E. Hart, and D. G. Stork, *Pattern Classification*, 2nd ed. Hoboken, NJ, USA: Wiley, 2012.
- [26] K. Simonyan and A. Zisserman. (2014). "Very deep convolutional networks for large-scale image recognition." [Online]. Available: <https://arxiv.org/abs/1409.1556>
- [27] J. T. Springenberg, A. Dosovitskiy, T. Brox, and M. A. Riedmiller. (2014). "Striving for simplicity: The all convolutional net." [Online]. Available: <https://arxiv.org/abs/1412.6806>
- [28] G. Larsson, M. Maire, and G. Shakhnarovich. (2016). "FractalNet: Ultra-deep neural networks without residuals." [Online]. Available: <https://arxiv.org/abs/1605.07648>
- [29] K. He and J. Sun, "Convolutional neural networks at constrained time cost," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2015, pp. 5353–5360.
- [30] V. Nair and G. E. Hinton, "Rectified linear units improve restricted boltzmann machines," in *Proc. 27th Int. Conf. Mach. Learn.*, 2010, pp. 1–8. [Online]. Available: <http://www.icml2010.org/papers/432.pdf>
- [31] A. Krizhevsky, "Learning multiple layers of features from tiny images," M.S. thesis, Dept. Comput. Sci., Univ. Toronto, Toronto, ON, Canada, 2009.
- [32] A. Novikov, D. Podoprikin, A. Osokin, and D. P. Vetrov, "Tensorizing neural networks," in *Proc. Adv. Neural Inf. Process. Syst.*, 2015, pp. 442–450.
- [33] W. Chen, J. Wilson, S. Tyree, K. Q. Weinberger, and Y. Chen, "Compressing neural networks with the hashing trick," in *Proc. Int. Conf. Mach. Learn.*, 2015, pp. 1–10.
- [34] M.-A. Côté and H. Larochelle, "An infinite restricted Boltzmann machine," *Neural Comput.*, vol. 28, no. 7, pp. 1265–1288, 2016.
- [35] C. Cortes, X. Gonzalvo, V. Kuznetsov, M. Mohri, and S. Yang. (2016). "AdaNet: Adaptive structural learning of artificial neural networks." [Online]. Available: <https://arxiv.org/abs/1607.01097>



Enrique S. Marquez received the B.S. degree in electrical engineering from Universidad Rafael Urdaneta, Maracaibo, Venezuela, in 2013, and the M.Sc. degree in artificial intelligence from the University of Southampton, Southampton, U.K., in 2015, where he is currently pursuing the Ph.D. degree in computer science.

His current research interests include machine learning, computer vision, and image processing.



Jonathon S. Hare received the B.Eng. degree in aerospace engineering and the Ph.D. degree in computer science from the University of Southampton.

He is currently a Lecturer in computer science with the University of Southampton, Southampton, U.K. His current research interests include multimedia data mining, analysis, and retrieval. These research areas are at the convergence of machine learning and computer vision, but also encompass other modalities of data. The long-term goal of his research is to innovate techniques that can allow machines to learn to understand the information conveyed by multimedia data and use that information to fulfill the information needs of humans.

Mahesan Niranjan held academic positions with the University of Sheffield, Sheffield, U.K., from 1998 to 2007, where he was the Head of the Department of Computer Science and the Dean of the Faculty of Engineering, and the University of Cambridge, Cambridge, U.K., from 1990 to 1997. He is currently a Professor of electronics and computer science with the University of Southampton, Southampton, U.K. His current research interest includes machine learning and made contributions to the algorithmic and applied aspects of the machine learning.