



How Deep Should be the Depth of Convolutional Neural Networks: a Backyard Dog Case Study

Alexander N. Gorban^{1,2} · Evgeny M. Mirkes^{1,2} · Ivan Y. Tyukin^{1,2,3}

Received: 2 October 2018 / Accepted: 3 July 2019
© The Author(s) 2019

Abstract

The work concerns the problem of reducing a pre-trained deep neuronal network to a smaller network, with just few layers, whilst retaining the network's functionality on a given task. In this particular case study, we are focusing on the networks developed for the purposes of face recognition. The proposed approach is motivated by the observation that the aim to deliver the highest accuracy possible in the broadest range of operational conditions, which many deep neural networks models strive to achieve, may not necessarily be always needed, desired or even achievable due to the lack of data or technical constraints. In relation to the face recognition problem, we formulated an example of such a use case, the 'backyard dog' problem. The 'backyard dog', implemented by a lean network, should correctly identify members from a limited group of individuals, a 'family', and should distinguish between them. At the same time, the network must produce an alarm to an image of an individual who is not in a member of the family, i.e. a 'stranger'. To produce such a lean network, we propose a network shallowing algorithm. The algorithm takes an existing deep learning model on its input and outputs a shallowed version of the model. The algorithm is non-iterative and is based on the advanced supervised principal component analysis. Performance of the algorithm is assessed in exhaustive numerical experiments. Our experiments revealed that in the above use case, the 'backyard dog' problem, the method is capable of drastically reducing the depth of deep learning neural networks, albeit at the cost of mild performance deterioration. In this work, we proposed a simple non-iterative method for shallowing down pre-trained deep convolutional networks. The method is generic in the sense that it applies to a broad class of feed-forward networks, and is based on the advanced supervise principal component analysis. The method enables generation of families of smaller-size shallower specialized networks tuned for specific operational conditions and tasks from a single larger and more universal legacy network.

Keywords Non-iterative learning · Principal component analysis · Convolutional neural networks

✉ Ivan Y. Tyukin
I.Tyukin@le.ac.uk

Alexander N. Gorban
a.n.gorban@le.ac.uk

Evgeny M. Mirkes
em322@le.ac.uk

- ¹ University of Leicester, Leicester LE1 7RH, UK
- ² Lobachevsky State University of Nizhny Novgorod, Prospekt Ganarina 23, Nizhny Novgorod, 603950, Russian Federation
- ³ Saint-Petersburg State Electrotechnical University (LETI), Prof. Popova 5, Saint-Petersburg, Russian Federation

Introduction

With the explosive pace of progress in computing, availability of cloud resources and open-source dedicated software frameworks, current artificial intelligence (AI) systems are now capable of spotting minute patterns in large data sets and may outperform humans and early-generation AIs in highly complicated cognitive tasks including object detection [1], medical diagnosis [2] and face and facial expression recognition [3, 4]. At the centre of these successes are deep neural networks and deep learning technology [5, 6].

Despite this, several fundamental challenges remain which constrain and impede further progress. In the

context of face recognition [4], these include the need for larger volumes of high-resolution and balanced training and validation data as well as the inevitable presence of hardware constraints limiting training and deployment of large models. Consequences of imbalanced training and testing data may have significant performance implications. At the same time, hardware limitations, such as memory constraints, restrict adoption, development and spread of technology. These challenges constitute fundamental obstacles for creation of universal data-driven AI systems, including for face recognition.

The challenge of overcoming hardware limitations whilst maintaining functionality of the underlying AI received significant attention in the literature. Heuristic definition of an efficient neural network was proposed in 1993: delivery of maximal performance (or skills) with minimal number of connections (parameters) [7]. Various algorithms of neural networks optimization were proposed in the beginning of the 1990s [8, 9]. MobileNet [10], SqueezeNet [11], DeepRebirth [12] and EfficientNets [13] are more recent examples of the approaches in this direction. Notwithstanding, however, the need for developing generic and flexible universal systems for a wide spectrum of tasks and conditions, there is a range of practical problems in which such universality may not be needed or required. These tasks may require smaller volumes of data and could be deployed on cheaper and accessible hardware. It is hence imperative that these tasks are identified and investigated, both computationally and analytically.

In this paper, we present and formally define such a task in the remit of face recognition: the ‘backyard dog’ problem. The task, on the one hand, appears to be a close relative of the standard face recognition problem. On the other, it is more relaxed which enables us to lift limitations associated with the availability of data and computational resources. For this task, we propose a technology and an algorithm for constructing a family of the ‘backyard dog’ networks derived from larger pre-trained legacy convolutional neural nets (CNN). The idea to exploit existing pre-trained networks is well known in the face recognition literature [14–18]. Our algorithm shares some similarity to [18] in that it exploits existing parts of the legacy system and uses them in a dedicated post-processing step. In our case, however, we apply these steps methodically across all layers; at the post-processing step, we employ advanced supervised principal component analysis (PCA) [19, 20] rather than conventional PCA, and do not use support vector machines.

Implementation of the technology and performance of the algorithm is illustrated with a particular network architecture, VGG net [15], and implemented on two computational platforms. The first platform was Raspberry Pi 3B with Broadcom BCM2387 chipset, 64-bit CPU

1.2 GHz Quad-Core ARM Cortex-A53 and 1 GiB memory with OS Raspbian Jessie. We will refer to it as ‘Pi’. The second platform was HP EliteBook laptop with Intel Core i7-840QM (4 x 1.86 GHz) CPU and 8 GiB of memory with OS Windows 7. We refer to this platform as ‘Laptop’. In view of Pi3 memory limitations (1 GiB), we required that the ‘backyard dog’ occupies no more than than 300 MiB. The overall workflow, however, is generic and should transfer well to other models and platforms.

The manuscript is organized as follows: in Section “[Preliminaries and Problem Formulation](#)”, we review the conventional face recognition problem, formulate the ‘backyard dog’ problem, assess several popular deep network architectures and select a test-bed architecture for implementation; Section “[The ‘backyard dog’ Generator](#)” describes the proposed shallowing technology for creation of the ‘backyard dog’ nets and illustrates it with an example; Section “[Conclusion](#)” concludes the paper.

Preliminaries and Problem Formulation

Face recognition is arguably among the hardest technical and computational problems. If posed as a conventional multi-class classification problem, it is ill-defined as acquiring samples from all classes, i.e. all identifies, is hardly possible. Therefore, state-of-the-art modern face recognition systems do not approach it as the multi-class classification problem. Not at least at the stage of deployment. These systems are often asked to answer another question: whether two given images correspond to the same person or not.

The common idea is to map these images into a ‘feature space’ with some metric (or a similarity measure) ρ . The system is then trained to ensure that if x and y are images corresponding to the same person then, for some $\varepsilon > 0$, $\rho(x, y) < \varepsilon$, and $\rho(x, y) > \varepsilon$ otherwise. At the decision stage, if $\rho(x, y) < \varepsilon$ then x, y represent the same person, and if $\rho(x, y) > \varepsilon$ then they belong to different identities. The problem with these generic systems is that validation and performance quantification for such systems is challenging; they must work well for all persons and images, including for identities these systems have never seen before.

It is thus hardly surprising that reports about performance of neural networks in face recognition tasks are often over-optimistic, with the accuracy of 98% and above [15–17] demonstrated on few benchmark sets. There is a mounting evidence that the training set bias, often present in face recognition datasets, leads to deteriorated performance in real-life applications [23]. If we use a human as a benchmark, trained experts make 20% mistakes on the faces they have never seen before [24]. Similar performance figures have been reported for modern face recognition systems when they assessed identities from populations that were

underrepresented in the training data [23]. Of course, we must always strive to achieve most ambitious goals, and the grand face recognition challenge is not an exception. Yet, in a broad range of practical situations, generality of the classical face recognition problem is not always needed or desired.

In what follows, we propose a relaxation of the face recognition problem that is significantly better defined and is closer to the standard multi-class problem with *known classes*. We call this problem the ‘backyard dog’ problem of which the specification is provided below.

The ‘backyard dog’ problem (Task) Consider a limited group of individuals, referred to as ‘family members’ (FM) or ‘friends’. Individuals who are not members of the family are referred to as ‘strangers’. A face recognition system, ‘the backyard dog’, should (i) separate images of friends from that of strangers and, at the same time (ii) should distinguish members of the family from each other (identity verification).

More formally, if q is an image of a person p , and Net is a ‘backyard dog’ net, then $Net(q)$ must return an identity class of q if $p \in FM$ and a label indicating the class of ‘strangers’ if $p \notin FM$.

The ‘backyard dog’ problem (Constraints) The ‘backyard dog’ must generate decisions within a given time frame on a given hardware and occupy no more than a given volume of RAM.

The difference between the ‘backyard dog’ problem and the traditional face recognition task is twofold. First, the ‘backyard dog’ should be able to reliably discriminate between a relatively small set of *known* identity classes (members of the family in the ‘backyard dog’ problem) as opposed to the challenge of reliable discrimination between pairs of images from a huge set of *unknown* identity classes (traditional face recognition setting). This is a significant relaxation as existing collections of training data used to develop models for face recognition (see Table 1) are several orders of magnitude smaller than 7.6 billion of the total world population [25]. In addition, the ‘backyard dog’ must separate a relatively small set of *known* friends from the huge but *unknown* set of potential strangers. The latter task is still challenging but its difficulty is largely reduced

relative to the original face recognition problem in that it is now a binary classification problem.

In the next sections, we will present a solution to the ‘backyard dog’ problem in which we will take advantage of the availability of a pre-trained deep legacy system. Before, however, presenting the solution lets us first select a candidate for a legacy system that would allow us to illustrate the concept better. For this purpose, below we review and assess some of the well-known existing system.

VGG

The Oxford Visual Geometry Group (and hence the name VGG) published their version of CNN for face recognition in [15]. We call this network VGGCNN [26]. The network was trained on a database containing facial images of 2622 different identities. Small modification of this network allows to compare two images and decide whether these two images correspond to the same person or not.

VGGCNN contains about 144M of weights. The recommended test procedure is as follows [15]:

1. Scale detected face to three sizes: 256, 384 and 512.
2. Crop a 224×224 fragment from each corner and from the centre of the scaled image.
3. Apply horizontal flip to crops.

Therefore, to test one face (one input image), one has to process 30 pre-processed images: $3(\text{scales}) \times 5(\text{crops}) \times 2(\text{flip}) = 30(\text{images})$.

Processing one image in the MatLab implementation [27] on our Laptop took approximately 0.7s. TensorFlow implementation [28] of the same required circa 7.3s.

FaceNet

Several CNNs with different architectures have been associated with the name [16]:

- NN1 with images 220×220 , 140M of weights and 1.6B FLOP,
- NN2 with images 224×224 , 7.5M of weights and 1.5B FLOP,

Table 1 Comparison of the datasets used to develop face recognition systems: (the table is presented in [15])

Dataset	Identities	Images	Link
LFW	5749	13,233	http://vis-www.cs.umass.edu/lfw/#download
WDRRef [21]	2995	99,773	N/A
CelebFaces [22]	10,177	202,599	N/A
VGG [15]	2622	2.6M	http://www.robots.ox.ac.uk/~vgg/data/vgg_face/
FaceBook [17]	4030	4.4M	N/A
Google [16]	8M	200M	N/A

Table 2 Memory requirements and computation resources needed: ‘Weights’ is the number of weights in the entire network in millions of weights; ‘Features’ is the maximal number of signals, in millions, which are passed from one layer to the other in a given network

Developer	Family name	Name	Weights (M)	Features (M)	FLOP (M)	Image size
VGG group	VGGCNN [15]	VGG16	144.0	6.4	15,475	224
Google	FaceNet [16]	NN1	140.0	1.2	1606	220
		NN2	7.5	2.0	1600	224
		NN3	7.5	NA	744	160
		NN4	7.5	NA	285	96
FaceBook	DeepFace [17]	DeepFace-align2D	118.0	0.8	805	152

- NN3 with images 160×160 , 7.5M of weights and 0.744B FLOP,
- NN4 with images 96×96 , 7.5M of weights and 0.285B FLOP.

Here, FLOP stays for floating point operations per image processing. The testing procedure for FaceNet uses one network evaluation per image.

DeepFace

FaceBook [17] proposed DeepFace architecture which, similarly to VGG face, is initially trained within a multi-class setting. At the evaluation stage, two replicas of the trained CNN assess a pair of images and produce their corresponding feature vectors. These are then passed into a separate network implementing the predicate ‘The same person/Different persons’.

Datasets

A comparison of the different datasets used to train the above networks is presented in Table 1. We can see that the

dataset used to develop VGG net is apparently the largest, except for the datasets used by Google, Facebook, or Baidu, which are not publicly available.

Comparison of VGGCNN, FaceNet and DeepFace

In addition to the training datasets, we have also compared the volumes of weights (in MiBs) and computational resources (in FLOPs) associated with each of the above networks. We did not evaluate their parallel/GPU-optimized implementations since our aim was to derive ‘backyard dog’ nets suitable for single-core implementations on the Pi platform. Results of this comparison are summarized in Tables 2 and 3. Distributions of weights, features and time needed to propagate an image through each network are shown in Figs. 1, 2, 3 and 4. Figures 1–4 also show that the interpretation of the notion of a ‘deep’ network varies for different teams: from 6 layers with weights in DeepFace to 16 such layers in VGG16.

According to Table 3, a C++ implementation for the Pi platform is comparable in terms of time with the TensorFlow (TF) implementation. Nevertheless, we note

Table 3 Computational time needed for passing one image through different networks

Developer	Family name	Name	Laptop ML	Laptop TF	Pi TF	Pi 1 core C++
VGG group	VGGCNN [15]	VGG16	0.695	4.723	75.301	65.909
Google	FaceNet [16]	NN1	0.072	0.490	7.815	6.840
		NN2	0.072	0.488	7.786	6.815
		NN3	0.033	0.227	3.620	3.169
		NN4	0.013	0.087	1.387	1.214
FaceBook	DeepFace [17]	DeepFace-align2D	0.036	0.246	3.917	3.429

For MatLab (ML) and TensorFlow (TF) realizations of VGGCNN on the Laptop platform, time was measured explicitly. All other values were estimated using FLOP data shown in Table 2 and taking VGGCNN data as a reference. Values for the Pi platform were estimated on the basis of explicit measurements for the reduced network (so that it fits into the system’s memory) and then scaled up proportionally




Fig. 1 Distribution of the numbers of weights, features and required FLOP along the VGG16 network from the first layer (left) to outputs (right)

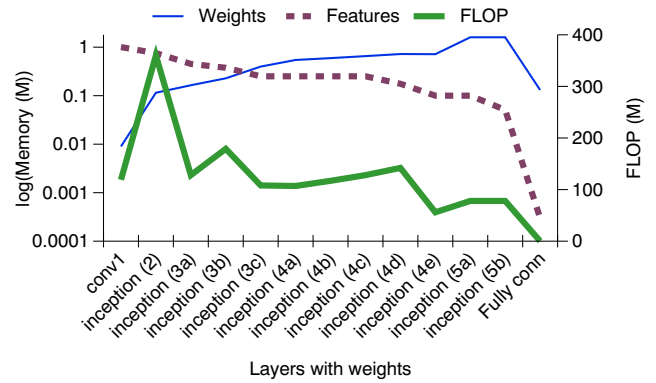


Fig. 3 Distribution of the numbers of weights, features and required FLOP along the NN2 network from the first layer (left) to outputs (right)

that we did not have control over the TF implementation in terms of enforcing the single-core operation. This may explain why single image processing times for the C++ and the TF implementations are so close.

In summary, we conclude that all these networks require at least 30 MiB of RAM for weights (7.5×4 MiB) and 3.2 MiB for features. Small networks (NN2-NN4) satisfy the imposed memory restrictions of 300 MiB. Large networks like VGG16, NN1 or DeepFace require more than 100 M of weights or 400 MiB and hence do not conform to this requirement. Time-wise, all candidate networks needed more than 1.2 s, with the VGGCNN requiring more than a minute on the Pi platform to process an image.

Having done this initial assessment, we therefore chose the largest and the slowest candidate as the legacy network. The task now is to produce a family of the ‘backyard dog’ networks from this legacy system which fit the imposed hardware constraints and, at the same time, deliver reasonable recognition accuracy. In the next section, we present a technology and an algorithm for creation of the ‘backyard dog’ networks from a given legacy net.




Fig. 2 Distribution of the numbers of weights, features and required FLOP along the NN1 network from the first layer (left) to outputs (right)

The ‘backyard dog’ Generator

Consider a general legacy network, and suppose that we have access to inputs and outputs for each layer of the network. Let the input to the first layer be an RGB image. One can now push this input through the first layer and generate this layer’s outputs. Output of the first layer becomes the first-layer features. For a multi-layer network, this process, if repeated throughout the entire network, will define features for each layer. At each layer, these features describe image characteristics that are relevant to the task which the network was trained on. As a general rule of thumb, as features of the deeper layers show higher degree of robustness. At the same time, this robustness comes at the price of increased memory and computational costs.

In our approach, we propose to seek a balance between the requirement of the task at hand, robustness (performance) and computational resources need. To achieve this balance, we suggest to assess suitability of the legacy system’s features layer by layer whereby determining the sufficient depth of the network and hence computational resources. The process is illustrated with Fig. 5. The ‘backyard dog’ net is a truncated legacy system whose outputs are fed into a post-processing routine.


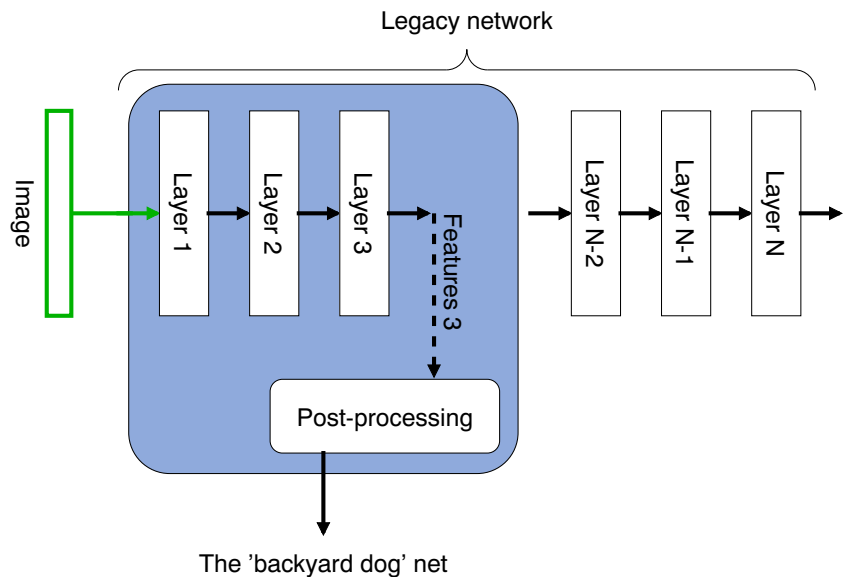


Fig. 4 Distribution of the numbers of weights, features and required FLOP along the DeepFace network from the first layer (left) to outputs (right)

Fig. 5 Construction of a ‘backyard dog’ network from a given legacy net



In principle, all layer types could be assessed. In practice, however, it may be beneficial to remove all fully connected layers from the legacy system first. This allows using image scaling as an additional hyper-parameter. This was the approach which we adopted here too.

The post-processing routine itself consisted of several stages:

- *Centralization*; subtraction of the mean vector calculated on the training set.
- *Spherical projection*; projection of the data onto a unit sphere centered at the origin (normalize each data vector to unit length).
- *Construction of new fully connected layer*; the output of this (linear in our case) layer is the output feature vector of the ‘backyard dog’.

Operational structure of the resulting network is shown in Fig. 6. Note that the first processing stage, centralization, can be implemented as a network layer subtracting a constant vector from the input. The second stage is a well-known L_2 normalization used, for example, in NN1 and NN2 [16]. As for the third stage, several approaches may exist. Here we will use advanced supervised PCA (cf. [18]). Details of the calculations used in relevant processing stages as well as interpretation of the ‘backyard dog’ net outputs are provided in the next section.




Fig. 6 Structure of created NN, abbreviations in brackets define notation of corresponding vectors

Interpretation of the ‘backyard dog’ Output Vector

Consider a set of identities, $P = \{p_1, \dots, p_n\}$, where n is the total number of persons in the database. A set of identities $FM = \{f_1, f_2, \dots, f_m\}$ forms a family (m is the number of FMs in the family). All identities, which are not elements of FM, are called ‘other persons’ or ‘strangers’. For each person f , $Im(f)$ is the set of images of this person, and $|Im(f)|$ is the total number of these images.

For an image q , we denote network output as $Out(q)$. Consider:

$$d(q) = \min_{f_i \in FM} \min_{r \in Im(f_i)} \|Out(q) - Out(r)\|. \quad (1)$$

Let $t >$ be a decision threshold. If $d(q) > t$ then the image q is interpreted as that of a non-family member (image of a ‘stranger’). If $d(q) \leq t$, then we interpret image q as that of FM f^* , where f^*

$$f^* = \arg \min_{f_i \in FM} \min_{r \in Im(f_i)} \|Out(q) - Out(r)\| \quad (2)$$

Three types of errors are considered:

MF: Misclassification of a FM. This error occurs when an image q belongs to a member of the set FM but $Out(q)$ is interpreted as ‘other person’ (a ‘stranger’).

MO: Misclassification of a ‘stranger’. This corresponds to a situation when an image q does not belong to any of identities from FM but $Out(q)$ is interpreted as FM.

MR: Misrecognition of a FM. This is an error when an image belongs to a member f_i of the set FM but $Out(q)$ is interpreted as an image of another FM.

Error rates are determined as the fractions of specific error types during testing (measured in %). The rate of **MF + MO** is the error rate of the ‘friend or foe’ task.

Construction of the ‘backyard dog’ Fully Connected (Linear) Layer

Interpretation rules above induce the following requirements for the new fully connected linear layer: we need to find an n -dimensional subspace S in the space of outputs such that the distance between projections onto S of the outputs corresponding to images of the same person is small and the distance between projections onto S of the outputs of images corresponding to different persons is relatively large. This problem has been considered and studied, for example, in [20, 29, 30]. Here we follow [19]. Recall that projection of the vector x onto the subspace defined by orthonormal vectors $\{v^i\}$ is Vx , where V is a matrix whose i th rows are v^i ($i = 1, \dots, n$). Select the target functional in the form:

$$D_C = D_B - \frac{\alpha}{k} \sum_{i=1}^k D_{W_i} \rightarrow \max, \tag{3}$$

where

k is the number of persons in the training set,
 D_B is the mean squared distance between projections of the network output vectors corresponding to different persons:

$$D_B = \frac{1}{\sum_{r=1}^{k-1} \sum_{s=r+1}^k |Im(p_r)| |Im(p_s)|} \times \sum_{r=1}^{k-1} \sum_{s=r+1}^k \sum_{x \in Im(p_r)} \sum_{y \in Im(p_s)} \|Vx - Vy\|^2, \tag{4}$$

D_{W_i} is the mean squared distance between projections of the network output vectors corresponding to person p_i :

$$D_{W_i} = \frac{1}{|Im(p_i)|(|Im(p_i)| - 1)} \sum_{x, y \in Im(p_i), x \neq y} \|Vx - Vy\|^2, \tag{5}$$

parameter α defines the relative cost for the output features corresponding to images of the same person being far apart.

The space of the n -dimensional linear subspaces of a finite-dimensional space (the Grassmannian manifold) is compact; therefore, the solution of Eq. 3 exists. The orthonormal basis of this space (the matrix V) is, by the definition, the set of the first n advanced supervised principal components (ASPC) [19]. They are the first n principal axis of the quadratic form defined from Eq. 3 [19, 20].

Training and Testing Protocol

In our case study, we used a database containing 25,402 images of 654 different identities [31] (38.84 images per person, on average). First, 327 identities were randomly selected from the database. These identities represented the set T of non-family members. Remaining 327 identities were used to generate sets of family members. We denote these identities as the set of family members candidates (FMC). Identities from this latter set with less than 10 images were removed from the set FMC and added to the set T of non-family members. From the set FMC, we randomly sampled 100 sets of 10 different identities, as examples of FM. We denote these sampled sets of identities as T_i , $i = 1, \dots, 100$. Elements of the set FMC which did not belong to any of the generated sets T_i were removed from the set FMC and added to the set T . As a result of this procedure, the set T contained 404 different identities.

For each truncated VGG16 network, and each image q of every identity in the training set T , we derived output vectors $VGG(q)$ and determined their mean vector $MVGG$

$$MVGG = \frac{1}{\sum_{f \in T} |Im(f)|} \sum_{f \in T} \sum_{q \in Im(f)} VGG(q). \tag{6}$$

This was used to construct the subtraction layer of which the output was defined as:

$$C(q) = VGG(q) - MVGG. \tag{7}$$

Each such vector $C(q)$ was then normalized to unit length.

Next, we determined ASPCs over the set of all vectors $C(q)$ associated with identities in the set T by solving (3). The value of alpha was varying in the interval [0.9, 2.3]. The value of t was chosen to minimize the rate of MF+MO error, for the given test set T_i , given value of α and the number of ASPCs. To determine optimal values of α and the number of ASPCs, we derived the mean values of MF, MO and MR across all T_i :

$$\begin{aligned} MF &= \frac{1}{100} \sum_{i=1}^{100} MF(T_i), \quad MO = \frac{1}{100} \sum_{i=1}^{100} MO(T_i), \\ MR &= \frac{1}{100} \sum_{i=1}^{100} MR(T_i) \end{aligned} \tag{8}$$

Table 4 Time, in seconds, spent on processing of a single image by different ‘backyard dog’ networks, columns T1 and T2 show outcomes of two identical tests executed at different times

Image size	Layers	ML		TF Laptop		TF Pi		C++	
		T1	T2	T1	T2	T1	T2	Laptop	Pi
224	37	0.67	0.72	7.35	7.05				
224	35	0.73	0.67						
224	31	0.62	0.66						
128	31	0.25	0.24						
96	31	0.19	0.21	0.96	0.95	17.08	17.31		
64	31	0.07	0.07	0.61	0.64	11.32	11.28		
96	24	0.12	0.13	0.59	0.43	7.44	8.91		
64	24	0.06	0.06	0.35	0.35	7.20	7.27	1.21	5.69
64	17							0.81	3.66
64	10							0.39	1.61
64	05							0.17	0.70

Table 5 Error rates for N05, N10, N17 and N24 without PCA improvement

Layers	MR	MF	MO	MF+MO
24	11.00	11.00	0.01	11.01
17	14.39	14.39	2.82	17.22
10	16.71	16.71	5.86	22.57
5	12.58	12.58	2.57	15.14

Error rates are evaluated as the maximal numbers of errors for 100 test sets (9)

Table 6 Error rates for N05, N10, N17 and N24 without PCA improvement

Layers	MR	MF	MO	MF+MO
24	4.16	4.13	1.09	5.22
17	7.69	7.65	1.75	9.39
10	10.94	10.82	3.64	14.46
5	6.58	6.52	2.01	8.53

Error rates are evaluated as the average numbers of errors for 100 randomly selected test sets (8)

Table 7 Error rates for networks with 5 and 17 layers and optimal number of ASPCs. Error rates are evaluated as the maximal numbers of errors for 100 randomly selected test sets (9)

Layers	MR	MF	MO	MF+MO
17	4.80	4.80	1.22	6.02
5	9.69	8.16	2.06	10.22

Table 8 Error rates for networks with 5 and 17 layers and optimal number of ASPCs, errors are evaluated as the average numbers of errors for 100 randomly selected test sets (8)

Layers	MR	MF	MO	MF+MO
17	2.50	2.46	0.81	3.27
5	4.39	4.30	1.48	5.78

as well as their maximal values

$$\begin{aligned} MF &= \max_i MF(T_i), MO = \max_i MO(T_i), \\ MR &= \max_i MR(T_i). \end{aligned} \quad (9)$$

For each of these performance metrics (Eqs. 8 and 9), we picked the number of ASPCs and the value of α which corresponded to the minimum of the sum MF+MO.

Results

Results of experiments are summarized in Tables 4, 5, 6, 7 and 8. Table 4 shows the amount of time each ‘backyard dog’ network required to process a single image. Tables 5–8 show performance of ‘backyard dog’ networks for varying depths (the number of layers). The best model for networks with 17 layers used 70 ASPCs, and the optimal network with 5 layers used 60 ASPCs.

The 5 layer network with 60 ASPCs processed a single 64×64 image in under 1 s on 1 core of Pi. It also demonstrated a reasonably good performance, with the MF+MO error rate below 6%. We note, however, that the reported performance levels in the ‘backyard dog’ problem are not to be confused with the system’s performance in more generic face recognition tasks. Note also that the maximal value of the MF+MO rate over 100 randomly selected sets T_i is 1.8 times higher than the average MF+MO rate for both 17 layer deep and 5 layer deep networks (with optimal number of ASPCs).

Conclusion

In this work, we proposed a simple non-iterative method for shallowing down legacy deep convolutional networks. The method is generic in the sense that it applies to a broad class of feed-forward networks, and is based on the ASPCA. We showed that, when applied to the state-of-the-art models developed for face recognition purposes, our approach generates a shallow network with reasonable performance in a specific task. The method enables one to produce of families of smaller-size shallower specialized networks tuned for specific operational conditions and tasks from a single larger and more universal legacy network.

The approach and technology were illustrated with a VGG-16 model. They will, however, apply to other models, including the popular MobileNet and SqueezeNet architectures. In this respect, our contribution is complementary to these works. Thanks to sufficiently large number of ASPCA projections used to produce ‘backyard dog’ net’s output, errors of the ‘backyard dog’ net may be reduced further using the error correction approach presented in [32–34]. Exploring this as well as testing the proposed approach on

other models, including MobileNet and SqueezeNet, will be the subject of our future work.

Acknowledgements We are grateful to Prof. Jeremy Levesley for numerous discussions and suggestions in the course of the project.

Funding This study was funded by by the Ministry of Science and Higher Education of Russian Federation (Project No. 14.Y26.31.0022) and Innovate UK Knowledge Transfer Partnership grants KTP009890 and KTP010522.

Compliance with Ethical Standards

Conflict of Interest The authors declare that they have no conflict of interest.

Ethical Approval This article does not contain any studies with human participants performed by any of the authors.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Krizhevsky A, Sutskever I, Hinton GE. Imagenet classification with deep convolutional neural networks. In: Advances in neural information processing systems; 2012. p. 1097–105.
- Huiying L, Tsui BY, Ni H, Valentim CCS, Baxter SL, Liu G, Cai W, et al. Evaluation and accurate diagnoses of pediatric diseases using artificial intelligence. *Nat Med.* 2019;25:433–8.
- Xiao S, Lv M. Facial expression recognition based on a hybrid model combining deep and shallow features. *Cognitive Computation.* 2019. <https://doi.org/10.1007/s12559-019-09654-y>.
- Ranjan R, Sankaranarayanan S, Bansal A, Bodla N, Chen J-C, Patel VM, Castillo CD, Chellappa R. Deep learning for understanding faces: machines may be just as good, or better, than humans. *IEEE Signal Process Mag.* 2018;35(1):66–83.
- Zhao Z-Q, Zheng P, Xu S-T, Wu X. Object detection with deep learning: a review. *IEEE transactions on neural networks and learning systems.* 2019.
- LeCun Y, Bengio Y, Hinton G. *Deep Learn Nat.* 2015; 521(7553):436–444.
- Gordienko P. Construction of efficient neural networks: algorithms and tests. In: Neural networks. IJCNN’93-Nagoya. Proceedings of 1993 International Joint Conference on 1993 Oct 25. IEEE; 1993. p. 313–6.
- Gorban AN. Training neural networks, USSR-USA JV ‘Paragraph’. 1990.
- Hassibi B, Stork DG, Wolff GJ. Optimal brain surgeon and general network pruning. In: IEEE International Conference on Neural Networks 1993. IEEE; 1993. p. 293–9.
- Howard AG, Zhu M, Chen B, Kalenichenko D, Wang W, Weyand T, Andreetto M, Adam H. Mobilenets: efficient convolutional neural networks for mobile vision applications. *arXiv:1704.04861.* 2017.
- Iandola FN, Han S, Moskewicz MW, Ashraf K, Dally WJ, Keutzer K. SqueezeNet: AlexNet-level accuracy with 50x fewer parameters and 0.5 MB model size. *arXiv:1602.07360.* 2016.

12. Li D, Wang X, Kong D. DeepRebirth: accelerating deep neural network execution on mobile devices. arXiv:1708.04728. 2017.
13. Mingxing T, Le QV. EfficientNet: rethinking model scaling for convolutional neural networks. arXiv:1905.11946. 2019.
14. Simonyan K, Zisserman A. Very deep convolutional networks for large-scale image recognition. In: International Conference on Learning Representations. 2015.
15. Parkhi OM, Vedaldi A, Zisserman A. Deep face recognition. In: Proceedings of the British Machine Vision Conference (BMVC). <http://www.robots.ox.ac.uk/vgg/publications/2015/Parkhi15/parkhi15.pdf>. 2015.
16. Schroff F, Kalenichenko D, Philbin J. Facenet: a unified embedding for face recognition and clustering. In: Proc. CVPR. 2015.
17. Taigman Y, Yang M, Ranzato M, Wolf L. Deep-face: closing the gap to human-level performance in face verification. In: Proc. CVPR. 2014.
18. Zhong G, Yan S, Huang K, Cai Y, Dong J. Reducing and stretching deep convolutional activation features for accurate image classification. *Cogn Comput*. 2018;10(1):179–86.
19. Mirkes EM, Gorban AN, Zinoviev A. Supervised PCA. <https://github.com/Mirkes/SupervisedPCA>. 2016.
20. Koren Y, Carmel L. Robust linear dimensionality reduction. *IEEE Trans Visual Comput Graph*. 2004;10(4):459–70. <https://doi.org/10.1109/TVCG.2004.17>.
21. Chen D, Cao X, Wang L, Wen F, Sun J. Bayesian face revisited: a joint formulation. In: Proc. ECCV. 2012; p. 566–79.
22. Sun Y, Wang X, Tang X. Deep learning face representation from predicting 10,000 classes. In: Proc. CVPR. 2014.
23. Lohr S. Face recognition is accurate, if you are a white guy. <https://www.nytimes.com/2018/02/09/technology/facial-recognition-race-artificial-intelligence.html>, The New York Times. 2018.
24. White D, Dunn JD, Schmid AC, Kemp RI. Error rates in users of automatic face recognition software. *PLOS One*. 2015;10(10):e0139827. <https://doi.org/10.1371/journal.pone.0139827>.
25. Population of the Earth. <http://www.worldometers.info/world-population/>.
26. Published VGG CNN <http://www.vlfeat.org/matconvnet/models/vgg-face.mat>.
27. MatConvNet <http://www.vlfeat.org/matconvnet>.
28. VGG in TensorFlow <https://www.cs.toronto.edu/frossard/post/vgg16/>.
29. Zinovyev AY. Visualisation of multidimensional data. Krasnoyarsk: Krasnoyarsk State Technocal University Press; 2000. In Russian.
30. Gorban AN, Zinovyev AY. Principal graphs and manifolds, chapter 2. Handbook of research on machine learning applications and trends: algorithms, methods, and techniques. In: Olivas ES et al., editors. Hershey: IGI Global; 2009. p. 28–59.
31. Gorban AN, Mirkes EM, Tyukin IY. Preprocessed database LITSO654 for face recognition testing <https://drive.google.com/drive/folders/10cu4u-31I24pKTOTIErjmie8gU-Z8biz?usp=sharing>.
32. Gorban AN, Golubkov A, Grechuk B, Mirkes EM, Tyukin I. Correction of AI systems by linear discriminants. *Probab Found Inf Sci*. 2018;466:303–22.
33. Tyukin I, Gorban AN, Green S, Prokhorov D. Fast construction of correcting ensembles for legacy artificial intelligence systems: algorithms and a case study. *Inf Sci*. 2019;485:230–47.
34. Gorban AN, Burton R, Romanenko I, Tyukin I. One-trial correction of legacy AI systems and stochastic separation theorems. *Inform Sci*. 2019;484:237–54.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.