

Face Model Compression by Distilling Knowledge from Neurons

Ping Luo^{1,3*}, Zhenyao Zhu^{1*}, Ziwei Liu¹, Xiaogang Wang^{2,3}, and Xiaoou Tang^{1,3}

¹Department of Information Engineering, The Chinese University of Hong Kong

²Department of Electronic Engineering, The Chinese University of Hong Kong

³Shenzhen Key Lab of Comp. Vis. & Pat. Rec., Shenzhen Institutes of Advanced Technology, CAS, China
{pluo,zz012,lz013,xtang}@ie.cuhk.edu.hk, {xgawang}@ee.cuhk.edu.hk

Abstract

The recent advanced face recognition systems were built on large Deep Neural Networks (DNNs) or their ensembles, which have millions of parameters. However, the expensive computation of DNNs make their deployment difficult on mobile and embedded devices. This work addresses model compression for face recognition, where the learned knowledge of a large teacher network or its ensemble is utilized as supervision to train a compact student network. Unlike previous works that represent the knowledge by the soften label probabilities, which are difficult to fit, we represent the knowledge by using the neurons at the higher hidden layer, which preserve as much information as the label probabilities, but are more compact. By leveraging the essential characteristics (domain knowledge) of the learned face representation, a neuron selection method is proposed to choose neurons that are most relevant to face recognition. Using the selected neurons as supervision to mimic the single networks of DeepID2+ and DeepID3, which are the state-of-the-art face recognition systems, a compact student with simple network structure achieves better verification accuracy on LFW than its teachers, respectively. When using an ensemble of DeepID2+ as teacher, a mimicked student is able to outperform it and achieves $51.6\times$ compression ratio and $90\times$ speed-up in inference, making this cumbersome model applicable on portable devices.

Introduction

As the emergence of big training data, Deep Neural Networks (DNNs) recently attained great breakthroughs in face recognition (Zhu et al. 2014a; 2013) and face verification (Taigman et al. 2014a; Zhu et al. 2014b; Sun, Wang, and Tang 2014; Schroff, Kalenichenko, and Philbin 2015) and become applicable in many commercial platforms, such as social networks, e-commerce, and search engines. To absorb massive supervision from big training data, existing works typically trained a large DNN or a DNN ensemble, where each DNN consists of millions of parameters. Nevertheless, as face recognition shifts toward mobile and embedded devices, large DNNs are computationally expensive, which prevents them from being deployed to these devices. It motivates research of using a small network to fit very large train-

ing data. This work addresses model compression of DNNs for face recognition, by incorporating *domain knowledge* of learning face representation.

There have been several attempts (Ba and Caruana 2014; Hinton, Vinyals, and Dean 2015; Romero et al. 2015) in literature to compress DNNs, so as to make their deployments easier, where a single network (*i.e.* a student) was trained by using the knowledge learned with a large DNN or a DNN ensemble (*i.e.* a teacher) as supervision. This knowledge can be simply represented as the probabilities of label predictions by employing the softmax function (Krizhevsky, Sutskever, and Hinton 2012). Compared with the original 1-of-K hard labels, the label probabilities encode richer relative similarities among training samples and can train a DNN more effectively. However, this representation loses much information because most of the probabilities are close to zeros after squashed by softmax. To overcome this problem, (Ba and Caruana 2014) represented the learned knowledge by using the logits, which are the values before softmax activation but zero-meaned, revealing the relationship between labels as well as the similarities among samples in the logit space. However, as these unconstrained values (*e.g.* the large negatives) may contain noisy information that overfits the training data, using them as supervision limits the generalization ability of the student. Recently, (Hinton, Vinyals, and Dean 2015) showed that both the label probabilities and zero-meaned logits are two extreme outputs of the softmax functions, where the temperature becomes one and positive infinity, respectively. To remove target noise, they empirically searched for a suitable temperature in the softmax function, until it produced soften probabilities that were able to disclose the similarity structure of data. As these soften target labels comprise much valuable information, a single student trained on them is able to mimic the performance of a cumbersome network ensemble. Despite the successes of (Hinton, Vinyals, and Dean 2015), our empirical results show that training on soft targets is difficult to converge when compressing DNNs for face recognition. Previous studies (Taigman et al. 2014b; Sun et al. 2014; Schroff, Kalenichenko, and Philbin 2015) have shown that the face representation learned from classifying larger amount of identities in the training data (*e.g.* 250 thousand in (Taigman et al. 2014b)) may have better generalization capacity. In face recognition, it seems diffi-

* indicates co-first authors who contributed equally.

Copyright © 2016, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

cult to fit soft targets with high dimensionality, which makes convergence slow.

In this work, we show that instead of using soft targets in the output layer, the knowledge of the teacher can also be obtained from the neurons in the top hidden layer, which preserve as much information as the soft targets (as the soft targets are predicted from these neurons) but are more compact, *e.g.* 512 versus 12,994 according to the net structure in (Sun, Wang, and Tang 2014). As these neurons may contain noise or information not relevant to face recognition, they are further selected according to the usefulness of knowledge captured by them. In particular, the selection is motivated by three original *observations* (domain knowledge) of face representation disclosed in this work, which are naturally generalized to all DNNs trained by distinguishing massive identities. (1) Deeply learned face representation by the face recognition task is a *distributed representation over face attributes*, including the identity-related attributes (IA), such as gender, race, and shapes of facial components, as well as the identity non-related attributes (NA), such as expression, lighting, and photo quality. This observation implies that each attribute concept is explained by having some neurons being activated while each neuron is involved in representing more than one attribute, although attribute labels are not provided during training. (2) However, a certain amount of neurons are selective to NA or both NA and IA, implying that the distributed representation is *neither invariant nor completely factorized*, because attributes in NA are variations that should be removed in face recognition, whereas these two factors (NA and IA) are presented and coupled in some neurons. (3) Furthermore, a small amount of neurons are inhibitive to all attributes and server as noise. With these observations, we cast neuron selection as inference on a fully-connected graph, where each node represents attribute-selectiveness of neuron and each edge represents correlation between neurons. An efficient mean field algorithm enables us to select neurons that are *more selective or discriminative to IA*, but *less correlated* with each other. As a result, the features of the selected neurons are able to maintain the inter-personal discriminativeness (*i.e.* distributed and factorized to explain IA), while reducing intra-personal variations (*i.e.* invariant to NA). We employ the features after neuron selection as regression targets to train the student.

To evaluate neuron selection, we employ DeepID2+ (Sun, Wang, and Tang 2014) as a teacher (T1), which achieved state-of-the-art performance on LFW benchmark. This work is chosen as an example because it successfully incorporated multiple complex components for face recognition, such as local convolution, ranking loss function, deeply supervised learning, and model ensemble. The effectiveness of all these components in face recognition have been validated by many existing works (Schroff, Kalenichenko, and Philbin 2015; Taigman et al. 2014b). Evaluating neuron selection on it demonstrates its capacity and generalization ability on mimicking functions induced by different learning strategies in face recognition. With neuron selection, a student with simple network structure is able to outperform a single network of T1 or its ensemble. Interestingly, this simple student generalizes well to mimic a deeper teacher (T2), DeepID3 (Sun

et al. 2015), which is a recent extension of DeepID2+.

Three main *contributions* of this work are summarized as below. (1) We demonstrate that more compact supervision converge more efficiently, when compressing DNNs for face recognition. Soft targets are difficult to fit because of high dimensionality. Instead, neurons in the top hidden layers are proper supervision, as they capture as much information as soft targets but more compact. (2) Three valuable observations are disclosed from the deeply learned face representation, identifying the usefulness of knowledge captured in these neurons. These observations are naturally generalized to all DNNs trained on face images. (3) With these observations, an efficient neuron selection method is proposed for model compression and its effectiveness is validated on T1 and T2.

Face Model Compression Training Student via Neuron Selection

The merit behind our method is to select informative neurons in the top hidden layer of a teacher, and adopt the features (responses) of the chosen neurons as supervision to train a student, mimicking the teacher’s feature space. We formulate the objective function of model compression as a regression problem given a training set $\mathcal{D} = \{\mathbf{I}_i, \mathbf{f}_i\}_{i=1}^M$,

$$L(\mathcal{D}) = \frac{1}{2M} \sum_{i=1}^M \|\mathbf{f}_i - g(\mathbf{I}_i; \mathbf{W})\|_2^2, \quad (1)$$

where \mathbf{I}_i and \mathbf{f}_i represent the i -th face image and its corresponding selected features, respectively. \mathbf{f}_i is obtained from a well training large DNN, which is the teacher. When dealing with an ensemble of DNNs, \mathbf{f}_i is selected from the top layers of all the DNNs. \mathbf{W} denotes a set of parameters of the student network and $g(\cdot)$ indicates a non-linear transformation from the input image to the features. Eqn.(1) is the objective function of training student network, which can be optimized by the stochastic gradient descent with standard back-propagation (BP) (Krizhevsky, Sutskever, and Hinton 2012).

Here, we introduce how to obtain the features \mathbf{f}_i in Eqn.(1) by selecting informative neurons. We formulate neuron selection as an inference problem on a fully-connected graph, where each node represents a neuron and each edge represents the correlation between a pair of neurons. Each node is associated with a binary latent variable, $y_i \in \{0, 1\}$, indicating whether neuron i has been chosen. Given a set of variables of N neurons, $\mathbf{y} = \{y_i\}_{i=1}^N$, the graph is optimized by minimizing the following energy function

$$E(\mathbf{y}) = \sum_{i=1}^N \Phi(y_i) + \lambda \sum_{i=1}^N \sum_{j=1, j \neq i}^N \Psi(y_i, y_j), \quad (2)$$

where $\Phi(y_i)$ and $\Psi(y_i, y_j)$ denote the unary and pairwise costs of selecting neuron i and both neurons i and j , respectively. λ is a constant weight. The first cost function is defined as $\Phi(y_i) = f(\mathbf{x}_i)$, where $f(\cdot)$ is a penalty function and \mathbf{x}_i is a vector measuring the attribute discriminativeness of neuron i . The second one measures the similarity between neurons, *penalizing* large correlation between




Figure 1: Attribute classification accuracies of single neurons in T1 are compared with the accuracies of single features from HOG and LBP.




Figure 2: (a) visualizes several neurons in the top hidden layer of T1, where the top one or two most dominative attributes for each neuron are outlined in yellow. For each neurons, images with low, medium, and high responses are grouped and averaged for visualization. (b) shows four different patterns of neuron activations. From left to right: neurons that are discriminative or selective to IA, NA, NA+IA, and the inhibitive neurons, respectively. Larger intensity indicates higher classification accuracy.

them, *i.e.* $\Psi(y_i, y_j) = \exp\{-\frac{1}{2}\|\mathbf{x}_i - \mathbf{x}_j\|_2^2\}$. These two terms demonstrate that we select neurons, which are more discriminative but less correlated. The representation of \mathbf{x} is discussed in the next section. The graphical model in Eqn.(2) can be solved by using the mean field algorithm (Jordan et al. 1999).

Attribute Discriminativeness of Neurons

We determine the values of \mathbf{x}_i for each neuron i , according to its selectiveness with respect to face attributes, which are implicitly captured by the learned face representation. To show this, we take the features of T1 (outputs of top hidden layer) as an example. Nevertheless, the following observations are naturally generalized to all DNNs trained to distinguish faces.

- Firstly, deeply learned face representation is a *distributed representation over face attributes*. This observation is inspired by (Sun, Wang, and Tang 2014), which showed that specific neurons at the higher layer are discriminative to 13 binary attributes that are closely related to face identity, such as gender and race. To further characterize the features, we employ the CelebA dataset (Liu et al. 2015) as a validation set, which contains 20 thousand face images and each image is annotated with 40 attributes. These attributes can mostly describe a face image as suggested by (Kumar et al. 2009). As shown in Fig.1, they include both identity-related attributes (IA) and non-related attributes (NA), plotted in black and gray respectively. We define each element j in vector \mathbf{x}_i as the mean classification accuracy of the j -th attribute, *i.e.* $\forall \mathbf{x}_i \in \mathbb{R}^{1 \times 40}$ and $\mathbf{x}_{i(j)} = \frac{TP_j + TN_j}{2}$, where

TP_j and TN_j represent the true positive and negative rates of attribute j , respectively. Fig.1 compares the maximum classification accuracy for each attribute achieved by using the neurons in T1 with the hand-crafted face descriptors like LBP and HOG¹, showing that deep features are automatically learned to distinguish *not only the attributes of IA but also NA*, although these concepts are not provided during training. However, LBP and HOG are not selective to neither of them. In this case, each identity is essentially represented by having a few neurons turned on (*i.e.* a distributed representation over attributes), while each neuron may be involved in representing more than one attribute as shown in Fig.2 (a), where several neurons are visualized by averaging face images of high, medium, and low responses.

- Secondly, this distributed representation is *neither invariant nor well factorized*, implying that some neurons are selective to NA or both IA and NA. As shown in Fig.2 (b), we randomly choose 80 neurons and partition them into four groups referring to their attribute-selectiveness, where larger intensity indicates higher accuracy. Most of the neurons are selective to IA, because they are trained to classify identities, but a quarter of neurons are sensitive to NA or IA+NA, implying that they are over-fitting the training data, since the attributes in NA such as ‘smiling’, ‘earrings’, and ‘lipstick’ belong to intra-personal variations to be removed in face recognition, while some neurons are not able to disen-

¹Each one-dimensional single feature of LBP or HOG is considered as an attribute classifier. The highest accuracy for each attribute achieved with a best single feature is reported.

(a) Teacher-1 (Sun, Wang, and Tang 2014) (T1): $55 \times 47 \times 3$ input image; $1 \times 12K$ output labels								
	1	2*	3	4*	5	6*	7	8*
<i>layer</i>	conv	max	conv	max	lconv	max	lconv	$2 \times$ fc
<i>neuron</i>	relu	—	relu	—	relu	—	relu	relu
<i>filter-stride</i>	4-1	2-2	3-1	2-2	3-1	2-2	2-1	—
<i>#channel</i>	128	128	128	128	128	128	128	1
<i>size</i>	52×44	26×22	24×20	12×10	10×8	5×4	4×3	$512, 12K$
<i>#param</i>	6K	37M, 6M	262K	8M, 6M	12M	1M, 6M	2M	1M, 6M

(b) Teacher-2 (Sun et al. 2015) (T2): $112 \times 96 \times 3$ input image; $1 \times 12K$ output labels										
	1	2	3	4*	5	6*	7	8*	9	10*
$2 \times$ conv	max	$2 \times$ conv	max	$2 \times$ conv	max	$2 \times$ conv	max	$2 \times$ lconv	$2 \times$ fc	
relu	—	relu	—	relu	—	relu	—	relu	relu	
3-1	2-2	3-1	2-2	3-1	2-2	3-1	2-2	3-1	—	
64	96	96	96	192	192	256	256	256	1	
112×96	56×48	56×48	28×24	28×24	14×12	14×12	7×6	$5 \times 4, 3 \times 2$	$512, 12K$	
39K	0	138K	17M, 3M	498K	8M, 3M	1M	3M, 3M	12M, 4M	1M, 6M	

(c) Student (S): $55 \times 47 \times 3$ input image; $1 \times N$ output								
	1	2	3	4	5	6	7	8*
<i>layer</i>	conv	max	conv	max	conv	max	fc	$2 \times$ fc
<i>neuron</i>	relu	—	relu	—	relu	—	relu	relu
<i>filter-stride</i>	4-1	2-2	3-1	2-2	3-1	2-2	—	—
<i>#channel</i>	128	128	128	128	128	128	1	1
<i>size</i>	52×44	26×22	24×20	12×10	10×8	5×4	500	$512, N$
<i>#param</i>	6K	0	262K	0	262K	0	1M	256K, $512 \times N$

Table 1: Comparisons among the network architectures of T1, T2, and the student model. Each table contains seven rows, representing the ‘type of layer’, ‘type of neuron’, ‘size of filter’–‘stride’, ‘number of channels’, ‘size of response map’, and ‘number of parameters’, respectively. Furthermore, ‘conv’, ‘lconv’, ‘max’, and ‘fc’ represent the convolution, local convolution, max pooling, and fully-connected layers respectively, while ‘relu’ indicates the rectified linear unit (Nair and Hinton 2010). For simplicity, thousand and million are denoted as ‘K’ and ‘M’.

tangle these factors.

- Thirdly, as illustrated at the rightmost side of Fig.2 (b), a small amount of neurons are inhibitive to all attributes, capturing no knowledge related to face and serving as noise.

As a good face representation should be both invariant and factorized to explain identity-related concepts, we select neurons discriminative to IA. To this end, the unary term can be written as $f(\mathbf{x}_i) = \frac{\max\{\mathbf{x}_{i(j)}\}_{\forall j \in \text{NA}} - \text{avg}\{\mathbf{x}_{i(j)}\}_{\forall j \in \text{NA}}}{\max\{\mathbf{x}_{i(j)}\}_{\forall j \in \text{IA}} - \text{avg}\{\mathbf{x}_{i(j)}\}_{\forall j \in \text{IA}}}$, where $\max\{\cdot\}$ and $\text{avg}\{\cdot\}$ look for the maximum and averaged values, respectively. If a neuron is more selective to NA compared to IA, $f(\cdot)$ produces large penalty, implying neurons discriminative to IA are more likely to be selected. Furthermore, the chosen neurons should have small correlations so as to explain different concepts. This constraint is modeled by the similarity between neurons as defined in the previous section. With the above definitions, we are able to select neurons by solving Eqn.(2).

Network Structures of Teachers and Student

This section introduces the structures of T1, T2, and a simple structure to mimic them.

Teachers The architectures of T1 and T2 are summarized in Table 1 (a) and (b) respectively, where the first two rows represent the types of layer and neuron, while ‘x-y’ in the third row represents the filter size and the stride of convolution. The last three rows represent number of channels, size of response maps, and number of parameters, respectively. As listed in Table 1 (a), T1 learns 512-dimensions face fea-

tures by classifying 12K identities with images of $55 \times 47 \times 3$ as input. It contains four convolutional layers, three max-pooling layers, and two fully-connected layers. These layers can be partitioned into eight groups, each of which covers one or more homogenous layers. The superscript (‘*’) over the group index indicates supervisory signals are propagated to this group. For instance, the second pooling layer is also connected to two ‘fc’ layers, which have the same hyper-parameters as group-8, leading to two large weight matrixes of $128 \times 26 \times 22 \times 512 \approx 37M$ and $512 \times 12K \approx 6M$ parameters, respectively. Similarly, the supervision is also propagated into group-4 and 6 respectively. In other words, T1 was trained in a deeply supervised manner, leading to more discriminative low- and middle-level representations. As a result, T1 has 85M parameters. T2 contains smaller number of parameters but deeper structure compared to T1. As listed in Table 1 (b), it has 62M parameters and 16-layers depth.

Student As shown in Table 1 (c), the structure of the student network (S) is simply derived from T1, where the local convolutional layers (‘lconv’) in group-5 and 7 are replaced by a convolutional layer and a fully-connected layer respectively, reducing the number of parameters by 11 times. All the supervision evaluated in our experiments are trained with S. For the soft targets (Hinton, Vinyals, and Dean 2015) and the logits (Ba and Caruana 2014), S is learned to minimize the cross-entropy and squared losses, respectively. Thus, the output dimension N equals 12K. For neuron selection, S predicts the features by minimizing Eqn.(1). In this case, N is typically smaller than 512, and therefore the student only

networks	#target length	AUC
T1 (single)	12,294	97.83
S-1-of-K	12,294	91.23
S-soft target ($t=1$)		94.42
S-soft target ($t=10$)		97.15
S-soft target ($t \rightarrow +\infty$)		96.75
S-neurons	512	97.73
S-selection	422	98.18

Figure 3: Mimicking a single network of T1.

contains about 2M parameters compared to the parameters of 85M in T1 and 62M in T2.

Experiments

Face Data For face model compression, we train all the students using data the same as (Sun, Wang, and Tang 2014), which combined two face databases for training, CelebFaces+ (Sun et al. 2014) and WDRRef (Chen et al. 2012), resulting in a training set of 290 thousands face images of 12,294 identities. In test, we evaluate all the models on LFW (Huang et al. 2007), which is the most well known benchmark for face recognition, containing 13,233 face images of 5,749 identities collected from the Internet. Note that the identities in training and test are exclusive. The face verification performance on LFW is reported as the Area under ROC curve (AUC) with respect to 3,000 positive and 3,000 negative face pairs.

For face verification, feature vectors in the top hidden layers are first extracted from a pair of face images and then the Euclidean distance² between them is calculated for face verification. Unlike (Taigman et al. 2014a; Sun, Wang, and Tang 2014) trained SVM or Joint Bayesian (Chen et al. 2012) for face verification, the Euclidean distance is used throughout the experiments to directly show the benefit from better supervision utilized to train students, other than strong classifiers with additional supervision.

Compressing a Single Model of T1 We train many students (S) with different targets to compress T1. Note that the structures of different S’s are the same except the last output layers. The architectures of S and T1 are given in Table 1. As shown in Fig.3, a student trained with knowledge distilled from selected neurons (‘S-selection’) achieves the best performance. It even outperforms its teacher T1, showing that selected neurons can preserve similarities among samples as well as remove noisy information in the features of T1. However, students supervised by the other targets have different losses in accuracy. Specifically, S is merely dropped by 0.68% compared to T1 when being trained with soft targets ($t = 10$), but dropped by 6.6% with hard labels (‘1-of-K’). Small networks cannot be well trained with hard labels. We examine different temperatures for soft targets. Note that when $t = 1$ and $t = +\infty$, soft targets turn into label probabilities (after softmax) and logits, respectively. The

²If the Euclidean distance between features extracted from a pair of face images exceeds a threshold, it is negative pair; otherwise, it is positive. The threshold is determined by cross-validation on LFW.

networks	#target length	AUC
T2 (single)	12,294	98.27
S-soft target ($t=1$)	12,294	91.88
S-soft target ($t=10$)		97.27
S-soft target ($t \rightarrow +\infty$)		96.77
S-neurons	512	97.90
S-selection	377	98.12
S-selection [†] (unsupervised)	377	98.37

Figure 4: Mimicking a single network of T2.

best performance of soft targets is achieved when $t = 10$. ‘S-neurons’ directly mimics the features of neurons in the top hidden layer of T1, and with 0.1% drop in the verification accuracy compared with T1. Neuron selection increases the accuracy by 0.35%. Furthermore, to verify the improvement is come from better targets, we train ‘S-1-of-K’ with attributes as additional supervision. The accuracy is 91%, showing that predicting identities and attributes jointly does not help face recognition, because attributes can be implicitly learned by classifying identities.

Fig.6 compares the training losses of T1 and different students. Several *valuable facts* can be observed. First, when using hard labels as targets, a larger network converges faster (e.g. comparing T1 and ‘S-1-of-K’), since it has larger fitting capability. Second, S’s trained with compact and informative targets converge faster than long targets, e.g. ‘S-neurons’ and ‘S-selection’ have 512 and 422 dimensional targets respectively, while ‘S-soft target’ has 12,294 dimensional targets. With these targets, a small student is able to learn the knowledge of a large teacher with a smaller amount of training time compared to that used to train the teacher. For example, ‘S-selection’ outperforms T1 with 42 times fewer parameters and learns at least 5 times faster. Third, convergence rate increases when the temperature of soft target increases. To the extreme, training with logits (i.e. $t = +\infty$) is easy, but training with 1-of-K labels is the most difficult, because the hard label vector has high dimensionality (e.g. 12K identities) and only one of its entries is non-zero. So the mimic network may produce wrong prediction at many iterations and converge slowly. In contrast, when temperature increases, the soft targets contain more non-zero values and become more informative.

Compressing a Deeper Teacher (T2) Fig.4 shows that the student (S) with shallow structure generalizes well to compress a deep teacher (T2), where T2 is $2 \times$ deeper than S. In this case, ‘S-selection[†]’ outperforms T2 by 0.1%. It is obtained by fine-tuning ‘S-selection’ with ten-folds cross-validation on LFW. This is done in an unsupervised manner without leveraging the identity labels of LFW, but using the selected features of T2. Without fine-tuning, the accuracy of ‘S-selection’ decreases by 0.15% compared to T2, indicating that deeper teacher is more difficult to fit. However, fine-tuning on more unsupervised data can improve performance. When compressing T2, S trained with neuron selection also outperforms those trained with the other supervision.

Compressing an Ensemble of T1 As listed in Fig.5, S is employed to compress an ensemble of six T1 networks, each

networks	#target length	AUC
T1 (an ensemble of 6 networks)	12,294	98.37
S-soft target (t=1)	12,294	91.88
S-soft target (t=5)		96.57
S-soft target (t=10)		96.07
S-soft target (t=15)		96.33
S-soft target (t $\rightarrow +\infty$)		96.65
S-neurons	3,072	98.07
S-selection	960	98.25
S-selection [†] (unsupervised)	960	98.43

Figure 5: Mimicking an ensemble of T1.




Figure 7: (a) Comparisons of running times (milliseconds per image). (b) Comparisons of number of parameters (millions).

of which was trained on different face regions, including two eyes, nose, two mouth corners, and the entire face region. This ensemble outperforms the best single network of T1 by 0.54%. For each student trained with soft targets, the supervision is obtained by averaging the soft targets of all the networks in this ensemble. When training ‘S-selection’ with the same data as T1 ensemble, its performance decreases by 0.12%. This result is similar to that of compressing T2, implying that not only deep teacher, an ensemble of teachers is also difficult to mimic. However, the performance can still be improved by unsupervised fine-tuning, *e.g.* ‘S-selection’ is increased by 0.18% after fine-tuning, outperforming T1 ensemble.

Complexity and Efficiency Fig.7 compare the efficiencies and complexities between T1, T2, T1 ensemble, and the student network (S). Efficiency is measured with implementation on a Intel Core 2.0GHz CPU. To simulate the environment of embedded or portable devices, the runtime is evaluated on CPU instead of GPU. As shown in Fig.7 (a), S achieves 90 \times speed-up compared to T1 ensemble. The model complexities are measured by the numbers of parameters during training and testing, respectively. *The former indicates learning capacity, while the latter indicates complexity in inference.* As shown in Fig.7 (b), if we compare S with T1 ensemble, S reduces the inference complexity by 51.6 \times and increases the performance by 0.06%, using a network structure with much smaller learning capacity, *i.e.* 1/255. In general, with neuron selection, the student S is able to outperform its corresponding teacher models by using much fewer parameters and process much faster. Specifically, it occupies 4 megabytes storage and processes face images with 250 frames per second, making T1 ensemble




Figure 6: Comparisons of convergence rates.

applicable on embedded or portable devices.

Conclusions and Discussions

This work demonstrates interesting results towards model compression for face recognition. In face recognition, both hard and soft labels are difficult to fit because of high dimensionality in the output layer, as well as the zero entries they contain. Instead, neurons in the top hidden layer are more suitable supervision because they capture as much information as soft targets, but are more compact. Experiments validate its effectiveness and its superior convergence rate. Valuable observations show that the deeply learned face representation is *neither invariant nor well factorized*. Therefore, employing all the features as targets is not a beneficial solution because they may contain noise or knowledge that is not relevant to face recognition. A neuron selection method is proposed to select neurons, so as to obtain a face representation that maintains the inter-personal discriminativeness, while reduces intra-personal variations. This is the goal of all the face recognition algorithms.

Acknowledgement

This work is partially supported by the National Natural Science Foundation of China (61503366, 91320101, 61472410), Guangdong Innovative Research Team Program (No.201001D0104648280), and Shenzhen Basic Research Program (KQCX2015033117354153, JCYJ20120903092050890, JCYJ20130402113127496).

References

- Ba, L. J., and Caruana, R. 2014. Do deep nets really need to be deep? In *NIPS*.
- Chen, D.; Cao, X.; Wang, L.; Wen, F.; and Sun, J. 2012. Bayesian face revisited: A joint formulation. In *ECCV*.
- Hinton, G. E.; Vinyals, O.; and Dean, J. 2015. Distilling the knowledge in a neural network. In *arXiv:1503.02531*.
- Huang, G. B.; Ramesh, M.; Berg, T.; and Learned-Miller, E. 2007. Labeled faces in the wild: A database for studying face recognition in unconstrained environments. In *University of Massachusetts, Amherst, Technical Report*.
- Jordan, M. I.; Ghahramani, Z.; Jaakkola, T. S.; and Saul, L. K. 1999. An introduction to variational methods for graphical models. In *Machine Learning, vol.37, no.2*.

Krizhevsky, A.; Sutskever, I.; and Hinton, G. E. 2012. Imagenet classification with deep convolutional neural networks. In *NIPS*.

Kumar, N.; Berg, A. C.; Belhumeur, P. N.; and Nayar, S. K. 2009. Attribute and simile classifiers for face verification. In *ICCV*.

Liu, Z.; Luo, P.; Wang, X.; and Tang, X. 2015. Deep learning face attributes in the wild. In *ICCV*.

Nair, V., and Hinton, G. E. 2010. Rectified linear units improve restricted boltzmann machines. In *ICML*.

Romero, A.; Ballas, N.; Kahou, S. E.; Chassang, A.; Gatta, C.; and Bengio, Y. 2015. Fitnets: Hints for thin deep nets. In *ICLR*.

Schroff, F.; Kalenichenko, D.; and Philbin, J. 2015. Facenet: A unified embedding for face recognition and clustering. In *arXiv:1503.03832*.

Sun, Y.; Chen, Y.; Wang, X.; and Tang, X. 2014. Deep learning face representation by joint identification-verification. In *NIPS*.

Sun, Y.; Liang, D.; Wang, X.; and Tang, X. 2015. Deepid3: Face recognition with very deep neural networks. In *arXiv:1502.00873*.

Sun, Y.; Wang, X.; and Tang, X. 2014. Deeply learned face representations are sparse, selective, and robust. In *arXiv:1412.1265*.

Taigman, Y.; Yang, M.; Ranzato, M.; and Wolf, L. 2014a. Deepface: Closing the gap to human-level performance in face verification. In *CVPR*.

Taigman, Y.; Yang, M.; Ranzato, M.; and Wolf, L. 2014b. Web-scale training for face identification. In *arXiv:1406.5266*.

Zhu, Z.; Luo, P.; Wang, X.; and Tang, X. 2013. Deep learning identity-preserving face space. In *ICCV*.

Zhu, Z.; Luo, P.; Wang, X.; and Tang, X. 2014a. Multi-view perceptron: a deep model for learning face identity and view representations. In *NIPS*.

Zhu, Z.; Luo, P.; Wang, X.; and Tang, X. 2014b. Recover canonical-view faces in the wild with deep neural networks. In *arXiv:1404.3543*.