Contents lists available at ScienceDirect





Pattern Recognition Letters

journal homepage: www.elsevier.com/locate/patrec

Large Age-Gap face verification by feature injection in deep networks

CrossMark

Simone Bianco

University of Milano-Bicocca, Viale Sarca 336, Ed. U14, Milan, 20126, Italy

ARTICLE INFO

Article history: Received 7 June 2016 Available online 14 March 2017

Keywords: Face verification Deep learning Cross-age face verification

ABSTRACT

This paper introduces a new method for face verification across large age gaps and also a dataset containing variations of age in the wild, the Large Age-Gap (LAG) dataset, with images ranging from child/young to adult/old. The proposed method exploits a deep convolutional neural network (DCNN) pre-trained for the face recognition task on a large dataset and then fine-tuned for the large age-gap face verification task. Fine-tuning is performed in a Siamese architecture using a contrastive loss function. A feature injection layer is introduced to boost verification accuracy, showing the ability of the DCNN to learn a similarity metric leveraging external features. Experimental results on the LAG dataset show that our method is able to outperform the face verification solutions in the state of the art considered.

© 2017 Elsevier B.V. All rights reserved.

1. Introduction

Face verification is an important topic in both computer vision, imaging and multimedia. Verification accuracy mainly depends on four elements: face pose, facial expression, illumination, and aging [17]. The greatest part of the works in the state of the art studied the face verification problem in constrained scenarios, controlling and fixing one or more of these four elements.

Recently many researchers achieved or even surpassed humanlevel performance [9,35] on face verification benchmark taken in unconstrained environments such as the Labeled Faces in the Wild dataset (LFW) [16]. These results have been made possible thanks to the improvement in facial landmark detection and to the increase of the computational power available to train deep models. However, the LFW dataset fixes the aging element: it contains large variations in pose, facial expression, and illumination, but contains very little variation in aging. As people grow, face appearance can be very different, which makes it difficult to recognize people across age. The problem is even harder when large age gaps are considered. In this work large age gap is interpreted in two ways: it refers both to the cases with extreme difference in age (e.g. young vs old) and to the cases with large difference in appearance due to the aging process (e.g. baby vs teenager/adult). Being able to recognize people across large age gaps could be beneficial in many applications: on photo sharing websites such as Facebook and Flickr that are providing services for over ten years; in all the personal photo management applications such as Google Photos and Apple Photo where albums can likely span decades; for the identification of long-lost and found persons.

http://dx.doi.org/10.1016/j.patrec.2017.03.006 0167-8655/© 2017 Elsevier B.V. All rights reserved. To address the problem of face verification across large age gaps, in this work a new approach is proposed. Differently from other approaches in the state of the art, the proposed method does not rely on parametric models nor tries to model age progression. The idea is to use deep learning to jointly learn face features that matching faces share, and a similarity metric on top of these features. This is done coupling two deep convolutional neural networks (DCNN) with shared parameters in a Siamese network [4,10] ended with a contrastive loss function. The discriminative power of the network is further improved including a feature injection layer, which fuses externally computed features with the activations of the deepest layers of the DCNN.

The contributions of this work are summarized as follows:

- A new large-scale Large Age-Gap (LAG) dataset is collected, that includes images in the wild of 1010 international celebrities spanning large age gaps.
- A new DCNN architecture is proposed, including a feature injection layer that fuses external features with the activations of the deepest DCNN layers.
- Extensive experiments are conducted on LAG and show that the proposed DCNN architecture can outperform state-of-theart methods. In particular, experimental results show that it is possible to increase the performance of a DCNN by injecting external features.

The remaining sections are organized as follows: Section 2 reviews the related works on face recognition, age-invariant face recognition and existing face datasets. Section 3 describes the proposed method, while Section 4 introduces the Large Age-Gap (LAG) dataset. Experiments are presented in Section 5. Finally, Section 6 draws the conclusions and discusses future works.

E-mail address: simone.bianco@disco.unimib.it

2.1. Deep feature fusion

The idea of deep feature fusion has been mainly explored in the video categorization task. One of the earliest work is from Simonyan and Zisserman [32] where they proposed a two-stream ConvNet architecture which incorporated a spatial and a temporal network. Perhaps the most similar work is [18] where multimodal video features are combined (e.g. frame-based features computed by a convolutional neural network, trajectory-based motion descriptors and audio descriptors). Wang et al. [38] integrate the advantages of hand-crafted and deep-learned features: they utilize deep architectures to learn multi-scale convolutional feature maps, and introduce the strategies of trajectory-constrained sampling and pooling to encode deep features into effective descriptors. Zha et al.[46] propose a late fusion approach between CNN features (taken at different layers) and Fisher Vectors [29]. The features are fused using an external classifier and thus not in an endto-end training, excluding the possibility of backward feedbacks on feature extraction. Ng et al. [23] investigated the combination of Long Short Term Memory (LSTM) networks [13] with optical flow. Park et al.[26] propose a multiplicative fusion method for combining multiple CNNs trained on different sources.

2.2. Face recognition

Face recognition has been investigated for a long time in many different works. One of the earliest works is that of Turk and Pentland where they introduced the idea of eigenface [36]. Ahonen et al. [1] explored the use of a texture descriptor, i.e. local binary pattern (LBP), for the face recognition task. Wright et al. [42] cast face recognition problem as one of classifying among multiple linear regression models via sparse signal representation, showing a high degree of robustness against face occlusions. Chen et al. [9] proposed a high dimensional version of LBP (HDLBP) and studied the performance of face feature as a function of dimensionality, showing that high dimensionality is critical to achieve high performance.

Recently there have been many works exploiting deep learning for face recognition. Results obtained by Taigman et al. [35] and by Sun et al. [33,34] using deep convolutional neural networks (DCNNs) reach or even surpass human-level performance on the widely used labeled face in the wild dataset (LFW) [16]. Their performance has been further improved by Schroff et al. [31] introducing FaceNet, that directly learns a mapping from face images to a compact Euclidean space where distances directly correspond to a measure of face similarity. Very recently Parkhi et al. [28] achieved even higher performance by using a very deep network that learns a face embedding using a triplet loss similar in spirit to that of Schroff et al. [31]. Although these methods achieve very high performance on face recognition, they do not work well when in presence of age variation, since this information is not used.

2.3. Age-Invariant face recognition

The largest part of existing works related to age in face image analysis focus on age estimation and simulation. Only recently researchers have started to work on cross age face recognition. Existing works can be grouped into generative and discriminative methods. Among the first group, some of the approaches build 2D [12] or 3D [27] aging models. These methods rely on parametric models and accurate age annotation or estimation, and thus do not work well in unconstrained scenarios. Wu et al. [43] propose a relative craniofacial growth model to model the facial shape change, which is based on the science of craniofacial anthropometry. Their method needs age information to predict the new shapes, limiting its applicability since this information is not always available.

Among the works based on a discriminative approach Li et al. [21] use multi-feature discriminant analysis (MFDA) to process in a unified framework the two local feature spaces generated by the two different local descriptors used, i.e. SIFT and multi-scale LBP. Gong et al. [14] proposed a method separating the HOG local feature descriptor into two latent factors using hidden factor analysis: an identity factor that is age-invariant and an age factor affected by the aging process. Chen et al. [6,7] use a data-driven method that leveraging a large-scale image dataset freely available on the internet as a reference set, encodes the low-level feature of a face image with an age-invariant reference space. Liu et al. [22] propose a generative-discriminative approach based on two modules: the aging pattern synthesis module and the aging face verification module. In the aging pattern synthesis module, an aging-aware denoising auto-encoder is used to synthesize the faces of all the four age groups considered. In the aging face verification module, parallel CNNs are trained based on the synthesized faces and the original faces to predict the verification score.

2.4. Face datasets

Existing face datasets can be divided into two main groups: the former consists of datasets acquired in controlled environments, the latter datasets in unconstrained environments. Most of the older datasets belong to the first group, such as FERET [30], Yale, and CMU PIE. The most popular dataset in uncontrolled environment is the LFW [16], with a total of 13,233 images of 5749 people extracted from news programs. Pubfig [20] has been collected with the aim of providing a larger number of images for each individual, and it contains 58,797 images of 200 identities. The largest dataset available is the CasiaWebFace dataset [44] with a total of 986,912 images of 10,575 people. All the above datasets can be used only for face recognition and verification tasks, since there is almost no age variation. Concerning age estimation and face recognition across age, the most used datasets are FGNet [11] and MORPH [30]. The former is composed of a total of 1002 images of 82 people with age range from 0 to 69 and an age gap up to 45 years. The latter contains 55,134 images of 13,618 people with age range from 16 to 77 and an age gap up to 5 years. Recently the CACD dataset has been collected [6,7] crawling the web suing as query 2000 celebrities names for a total of 163,446 images. For a subset of 200 identities images are manually checked and the noisy ones have been removed. Age ranges from 14 to 62 and age gap is up to 10 years. Very recently the CAFE dataset has been collected [22]. It is the first permitting a study on face verification with large age gaps. It is composed of 4659 images of 901 people and, due to the way images are collected it does not contain precise age information. A summary of the comparison between existing datasets is reported in Table 1.

3. The proposed method

Fig. 1 gives an overview of the proposed method. First face and landmark detection are performed on CASIA-WebFace and Large Age-Gap (LAG) database to localize and align each face to a reference position. Next, a DCNN is trained on the CASIAWebFace [44] for the face recognition task. The knowledge learned by the DCNN is then transferred [25] from the source task of face recognition to the target task of large age gap face verification. Knowledge transfer is performed by fine-tuning [2,45] the DCNN using a contrastive loss in a Siamese architecture in which pre-computed external features are injected in the fully connected layer.

-		-		-			
Dataset name	Year	# of images	# of people	# images/person	Age info.	Age gap	Publicly available
LFW [16]	2007	13,233	5749	2.3	No	-	Yes
Pubfig [20]	2009	58,797	200	293.9	No	-	Yes
Casia [44]	2014	986,912	10,575	93.3	No	-	Yes
FGNet [11]	2008	1002	82	12.2	Yes	0-45	Yes
MORPH [30]	2006	55,134	13,618	4.1	Yes	0-5	Yes
CACD [6,7]	2015	163,446	2000	81.7	Yes	0-10	Yes
CAFE [22]	2016	4,659	901	5.2	No	Large	Not yet
Ours (LAG)	this paper	3828	1010	3.8	No	Large	Yes*

The comparison between existing datasets for face verification/recognition.

* http://www.ivl.disco.unimib.it/activities/large-age-gap-face-verification/.



Fig. 1. Overview of the proposed method for large age-gap face verification.

3.1. Image preprocessing

For each image in the database, we first apply the widely used Viola–Jones face detector [37] to find the face region in the image. For each face, we then locate 68 different facial landmarks using a face alignment algorithm [19]. After landmark detection, we use the eye locations to align the face images. Images are first rotated so that the eyes are horizontally aligned, then scaled so that the distance between eyes is fixed, and finally cropped to a common size of 200×200 pixels.

3.2. Deep face feature representation

A DCNN is trained to learn a discriminative representation of faces. The chosen architecture is the AlexNet [47], but others could be used. The choice of this architecture was made since it has shown strong results in the literature and permits ease of comparison to other works. The DCNN is trained on the face recognition task on the CASIAWebFace. The dimensionality of the input layer is $200 \times 200 \times 1$ gray-scale images. The network includes 5 convolutional layers, 3 pooling layers and 3 fully connected layers. Each convolutional layer is followed by a rectified linear unit (ReLU). Two local normalization layers are added after each of the first two convolutional layers to mitigate the ef-

fect of illumination variations. Dropout is used to regularize all the fully connected layers due to the large number of parameters (fc6 = 4096, fc7 = 4097, and fc8 = 10575 equal to the number of different identities in CASIAWebFace). The features extracted from the second to last fully connected layer, i.e. fc7, are used for face representation after an L_2 -normalization step.

3.3. Feature injection

The L_2 -normalized fc7 features are given as input to a set of n face verification methods in the state of the art. Each of them provides as output a distance or confidence score that a pair of images belong to the same identity or not d_i , i = 1, ..., n, which are stacked to form the vector $\mathbf{d} = \{d_i\}_{i=1}^n$. A Siamese DCNN [4,10] is then fine-tuned on the LAG database starting from the net in Section 3.2 using a contrastive loss function. In addition to the DCNN features, a feature injection layer is added to fuse externally computed features with the activations of the deepest layers of the DCNN. In more details, the feature injection is performed in the form of concatenation of the external features \mathbf{d} with the fc6's activations, as represented in Fig. 2 for one side of the Siamese DCNN. The idea is that the DCNN jointly learns face features that matching faces share, and a similarity metric on top of these

Table 1



features

Fig. 2. Graphical representation of feature injection in the fully connected layer.



Fig. 3. Examples of face crops for matching pairs in the Large Age-Gap (LAG) dataset.

features also leveraging external features. At test time, the fc7's activations with injected features are computed for the pair of input images, then they are L_2 -normalized, and their Euclidean distance is thresholded to classify them as matching pairs or not.

4. Large Age-Gap dataset (LAG)

In order to be able to collect images from a large number of people, the Large Age-Gap dataset (LAG) is constructed with photos of celebrities. The meaning of large gap is twofold: from one side it refers to images with extreme difference in age, e.g. 0 to 80 years old; on the other side it refers to large difference in appearance due to the aging process: for example, 0 to 15 years old is a relatively small difference in age but has a large change in appearance. The LAG dataset covers both aspect.

Google Image Search is used to collect images by specifying to search face images. In order to collect celebrities images across different ages, we augment the celebrity name with adjectives such as "baby", "childhood", "adult", "now", "today" as keywords. Searches with celebrity names followed by the string "then and now" are also used. YouTube videos of collections of "then and now" celebrities have also been downloaded. Since the videos share a common and very simple structure, i.e. short sequences of the two still images (the younger and the older version) one next to the other, all the video frames are extracted and a simple shot detection algorithm is used to detect the duration of the different sequences. Then, the central frame of each shot is collected.

The next step is the duplicate image removal. To this end, a structural similarity measure [39] is used to compare all the possible couples of images from the same identity. Only the couples with a very high similarity are inspected, and if they are actually duplicates only the higher quality version is maintained. After removing duplicate images, we manually check the images and remove the noisy images in the dataset. The dataset contains 3828 images of 1010 celebrities after removing the noisy images. For

each identity at least one child/young image and one adult/old image are present. Starting from the collected images, a total of 5051 matching pairs has been generated. The same number of non matching pairs has been randomly generated. Table 1 shows the statistics of the dataset and comparison to other existing face datasets. Some examples of the face crops of the collected images are reported in Fig. 3.

5. Experimental results

In this section we evaluate the performance of large age-gap face verification. We compare our method with current state-of-the-art features and classifiers for general face verification, such as high dimensional local binary feature (HDLBP) [9], Other similarity metric learning approaches, such as sub-SML [5], One Shot Similarity Kernel [40], Cosine Similarity [24] and Joint Bayesian [8]. All the similarity learning methods have been trained using the L_2 -normalized *fc*7 features (extracted with the DCNN trained for the face recognition task on CASIAWebFace). We also compare our method with the Cross Age Reference Coding (CARC) [6,7] method, which is designed specifically for cross-age face verification. Since LAG dataset does not provide exact age information, the version without the temporal constraint is used here, i.e. CARC-NT.

To understand the contribution of the feature injection, a Siamese DCNN is fine-tuned on the LAG dataset excluding the injection layer. This also permits a direct comparison of the Siamese DCNN with the other similarity metric learning approaches considered. Besides feature injection, to have a benchmark, external features are also fused stacking [41] linear SVM classifiers as in Hassner et al. [15].

Experiments are made using a two-fold cross validation: the LAG dataset has been alphabetically sorted and subjects have been assigned alternately to the first or to the second fold. For each fold, the training set is augmented by considering the four combinations of horizontal flips and original images, as well as tiny amounts of

Average recognition accuracy comparison of our method and other baselines on the LAG dataset.

_					
	ID	Method	Features	Accuracy	Improvement
	1.	Euclidean distance	DCNN	0,5917	-0,1565
	2.	L2-norm & Euclidean Distance	DCNN	0,6498	-0,0984
	3.	L2-norm & Hellinger Distance	DCNN	0,6688	-0,0794
	4.	Similarity Metric Learning (sub-SML) [5]	DCNN	0,7243	-0,0239
	5.	OSS (one shot similarity) [40]	DCNN	0,6642	-0,0840
	6.	Cosine Similarity [24]	DCNN	0,6508	-0,0974
	7.	Joint Bayesian [8]	DCNN	0,6633	-0,0848
	8.	CARC-NT [6,7]	DCNN	0,7482	-,
	9.	HDLBP [9]	LBP	0,7153	-0,0329
	10.	Stacking (1 representation: DCNN)	1+2+3	0,6781	-0,0701
	11.	Stacking (1 representation: DCNN)	1+2+3+4	0,7441	-0,0041
	12.	Stacking (1 representation: DCNN)	1+2+3+4+5	0,7874	0,0392
	13.	Stacking (1 representation: DCNN)	1+2+3+4+5+7	0,7712	0,0231
	14.	Stacking (1 representation: DCNN)	1+2+3+4+5+8	0,7921	0,0439
	15.	Siamese DCNN	DCNN	0,7734	0,0252
	16.	Stacking (1 representation: DCNN)	1+2+3+4+5+8+15	0,8010	0,0528
	17.	Siamese DCNN + feature injection	DCNN with injected 1+2+3+4+5+8+15	0,8277	0,0795
	18.	Stacking (2 representations: DCNN + LBP)	1+2+3+4+5+8+9+15	0,8263	0,0781
	19.	Siamese DCNN + feaure injection	DCNN with injected 1+2+3+4+5+8+9+15	0,8495	0,1014



Fig. 4. ROC curves of the proposed method and other baselines on the LAG dataset.

jittering. Performance are reported as the verification accuracy in Table 2, and plotted as ROC curves in Fig. 4.

From the results reported in Table 2 it is possible to see that the use of the Siamese DCNN fine-tuned on the LAG dataset to learn a similarity metric is able to outperform all the single methods in the state of the art, outperforming the best one by almost 2.5%. Enabling the feature injection in the fine-tuning step permits an improvement over the best method in the state of the art of more than 10%. Feature combination stacking linear SVM classifiers performs 2.3% worse than our method. The experimental results show the effectiveness of both the use of a Siamese DCNN for similarity metric learning and the usefulness of the feature injection. An indirect comparison with the method by Liu et al. [22] can be done by using HDLBP[9] as a reference: on their CAFE dataset they reported an improvement over HDLBP of 2.5%. On our LAG dataset our methods outperforms HDLBP by 13.4%.

Table 3							
Confusion	matrix	of	our	method	on	the	LAG
dataset.							

	non-matching	matching
non-matching	0.8575	0.1425
matching	0.1584	0.8416

In the following we examine more in details the results of the best performing Siamese DCNN with feature injection reported in Table 2 (i.e. ID. 19): its confusion matrix is reported in Table 3. From the confusion matrix we can evince that our solution is almost equally able to identify non-matching pairs and matching pairs.

In Fig. 5 some examples of correctly identified matches among those on which our solution is very confident are reported. Some examples of correctly identified matches on which our solution is not so confident are reported in Fig. 6. Some examples of false positives, i.e. non-matching pairs classified as matching ones are reported in Fig. 7, while examples of misclassified matching pairs are reported in Fig. 8. From the examples reported it is possible to see that matching pairs on which our solution is very confident tend to have the same pose and expression, suggesting that more powerful pose-normalization methods could further improve the verification accuracy. A similar behavior can be observed on the false positives, where the similarity between specific facial features led the system to wrongly classify them as matching pairs (see for example the pointed chin in the 8th and 10th column, the eye expression in the 9th column, the expression and pose in the 7th column, or the expression and face shape in the 2nd column). For what concerns the examples of correctly identified matches on which our solution is not so confident, we can observe that in some cases there is a large change in appearance also due to a change in facial features e.g. the beard in the example in the 1st, 3rd and 9th column or eyeglasses in the 8th column. Some of the false negative examples are instead very difficult to classify since the subjects have had some sort of plastic surgery procedures (e.g. see Michael Jackson).

6. Conclusions

In this paper a new method for face verification across large age gaps is introduced along with a dataset containing variations of age in the wild. The proposed method exploits a deep



Fig. 5. Examples of correctly classified matching pairs on which our approach is very confident.



Fig. 6. Examples of correctly classified matching pairs on which our approach is not very confident.



Fig. 7. Examples of errors of our approach: non-matching pairs classified as matching ones (i.e. false positives).



Fig. 8. Examples of errors of our approach: matching pairs classified as non-matching ones (i.e. false negatives).

convolutional neural network (DCNN) trained in a Siamese architecture ended with a contrastive loss function. The network discriminative power is further improved including a feature injection layer, which injects externally computed features into the deepest layers of the DCNN. Experimental results on the Large Age-Gap (LAG) dataset show that the proposed approach is able to outperform the face verification methods in the state of the art considered.

As future work we plan to extend this research in different directions:

 Integration of the LAG dataset with the CAFE dataset [22] when it will be released, and further extension of the LAG dataset collecting more images for each identity.

- Investigation of the use of different face alignment algorithms such as Hassner et al. [15].
- Comparison and integration with face aging models [12].
- Addition of pre-processing steps such as illuminant compensation [3].
- Use of pre-classifiers to make verification task easier, e.g. age/gender/race classifier, expression classifier, etc.

Acknowledgments

The author gratefully acknowledges the support of NVIDIA Corporation with the donation of the Tesla K40 GPU used for this research.

References

- [1] T. Ahonen, A. Hadid, M. Pietikainen, Face description with local binary patterns: application to face recognition, Pattern Anal. Mach. Intell. IEEE Trans. 28 (12) (2006) 2037–2041.
- [2] Y. Bengio, et al., Deep learning of representations for unsupervised and transfer learning, in: ICML Unsupervised and Transfer Learning, 27, 2012, pp. 17–36.
- [3] S. Bianco, R. Schettini, Adaptive color constancy using faces, IEEE Trans. Pattern Anal. Mach. Intell. 36 (8) (2014) 1505–1518.
- [4] J. Bromley, J.W. Bentz, L. Bottou, I. Guyon, Y. LeCun, C. Moore, E. Säckinger, R. Shah, Signature verification using a siamese time delay neural network, Int. J. Pattern Recognit. Artif. Intell. 7 (04) (1993) 669–688.
- [5] Q. Cao, Y. Ying, P. Li, Similarity metric learning for face recognition, in: Computer Vision (ICCV), 2013 IEEE International Conference on, IEEE, 2013, pp. 2408–2415.
- [6] B.-C. Chen, C.-S. Chen, W. Hsu, Face recognition and retrieval using cross-age reference coding with cross-age celebrity dataset, Trans. Multimedia 17 (6) (2015) 804–815.
- [7] B.-C. Chen, C.-S. Chen, W.H. Hsu, Cross-age reference coding for age-invariant face recognition and retrieval, in: Computer Vision–ECCV 2014, Springer, 2014, pp. 768–783.
- [8] D. Chen, X. Cao, L. Wang, F. Wen, J. Sun, Bayesian face revisited: a joint formulation, in: Computer Vision-ECCV 2012, Springer, 2012, pp. 566–579.
- [9] D. Chen, X. Cao, F. Wen, J. Sun, Blessing of dimensionality: high-dimensional feature and its efficient compression for face verification, in: Computer Vision and Pattern Recognition (CVPR), 2013 IEEE Conference on, IEEE, 2013, pp. 3025–3032.
- [10] S. Chopra, R. Hadsell, Y. LeCun, Learning a similarity metric discriminatively, with application to face verification, in: Computer Vision and Pattern Recognition, 2005. CVPR 2005., 1, IEEE, 2005, pp. 539–546.
- [11] T. Cootes, A. Lanitis, The fg-net aging database, 2008,
- [12] X. Geng, Z.-H. Zhou, K. Smith-Miles, Automatic age estimation based on facial aging patterns, Pattern Anal. Mach. Intell. IEEE Trans. 29 (12) (2007) 2234–2240.
- [13] F.A. Gers, N.N. Schraudolph, J. Schmidhuber, Learning precise timing with lstm recurrent networks, J. Mach. Learn. Res. 3 (2003) 115–143.
- [14] D. Gong, Z. Li, D. Lin, J. Liu, X. Tang, Hidden factor analysis for age invariant face recognition, in: Proceedings of the IEEE International Conference on Computer Vision, 2013, pp. 2872–2879.
- [15] T. Hassner, S. Harel, E. Paz, R. Enbar, Effective face frontalization in unconstrained images, in: Computer Vision and Pattern Recognition (CVPR), 2015 IEEE Conference on, IEEE, 2015, pp. 4295–4304.
- [16] G.B. Huang, M. Ramesh, T. Berg, E. Learned-Miller, Labeled Faces in the Wild: A Database for Studying Face Recognition in Unconstrained Environments, Technical Report, 07–49, University of Massachusetts, Amherst, 2007.
- [17] A.K. Jain, B. Klare, U. Park, Face matching and retrieval in forensics applications, IEEE Multimedia 19 (1) (2012) 20.
- [18] Y.-G. Jiang, Z. Wu, J. Wang, X. Xue, S.-F. Chang, Exploiting feature and class relationships in video categorization with regularized deep neural networks, arXiv preprint arXiv:1502.07209(2015).
- [19] V. Kazemi, J. Sullivan, One millisecond face alignment with an ensemble of regression trees, in: Computer Vision and Pattern Recognition (CVPR), 2014 IEEE Conference on, IEEE, 2014, pp. 1867–1874.
- [20] N. Kumar, A.C. Berg, P.N. Belhumeur, S.K. Nayar, Attribute and simile classifiers for face verification, in: Computer Vision, 2009 IEEE 12th International Conference on, IEEE, 2009, pp. 365–372.
- [21] Z. Li, U. Park, A.K. Jain, A discriminative model for age invariant face recogni-
- tion, Inf. Forensics Secur. IEEE Trans. 6 (3) (2011) 1028–1037. [22] L. Liu, C. Xiong, H. Zhang, Z. Niu, M. Wang, S. Yan, Deep aging face verification
- with large gaps, Multimedia IEEE Trans. 18 (1) (2016) 64–75. [23] J.Y.-H. Ng, M. Hausknecht, S. Vijayanarasimhan, O. Vinyals, R. Monga, G.
- Toderici, Beyond short snippets: deep networks for video classification, arXiv preprint arXiv:1503.08909(2015).

- [24] H.V. Nguyen, L. Bai, Cosine similarity metric learning for face verification, in: Computer Vision–ACCV 2010, Springer, 2011, pp. 709–720.
- [25] S.J. Pan, Q. Yang, A survey on transfer learning, IEEE Trans. Knowl. Data Eng. 22 (10) (2010) 1345–1359.
- [26] E. Park, X. Han, T.L. Berg, A.C. Berg, Combining multiple sources of knowledge in deep cnns for action recognition, in: Winter Conference on Applications of Computer Vision (WACV), 2016.
- [27] U. Park, Y. Tong, A.K. Jain, Age-invariant face recognition, Pattern Anal. Mach. Intell IEEE Trans. 32 (5) (2010) 947–954.
- [28] O.M. Parkhi, A. Vedaldi, A. Zisserman, Deep face recognition, in: British Machine Vision Conference, 1, 2015, p. 6.
- [29] F. Perronnin, J. Sánchez, T. Mensink, Improving the Fisher Kernel for large-scale image classification, in: ECCV 2010, 2010, pp. 143–156.
- [30] K. Ricanek Jr, T. Tesafaye, Morph: a longitudinal image database of normal adult age-progression, in: Automatic Face and Gesture Recognition, 2006. FGR 2006. 7th International Conference on, IEEE, 2006, pp. 341–345.
- [31] F. Schroff, D. Kalenichenko, J. Philbin, Facenet: a unified embedding for face recognition and clustering, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2015, pp. 815–823.
- [32] K. Simonyan, A. Zisserman, Two-stream convolutional networks for action recognition in videos, in: Advances in Neural Information Processing Systems, 2014, pp. 568–576.
- [33] Y. Sun, Y. Chen, X. Wang, X. Tang, Deep learning face representation by joint identification-verification, in: Advances in Neural Information Processing Systems, 2014, pp. 1988–1996.
- [34] Y. Sun, X. Wang, X. Tang, Deep learning face representation from predicting 10,000 classes, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2014, pp. 1891–1898.
- [35] Y. Taigman, M. Yang, M. Ranzato, L. Wolf, Deepface: closing the gap to human-level performance in face verification, in: IEEE Conference on Computer Vision and Pattern Recognition, 2014, pp. 1701–1708.
- [36] M.A. Turk, A.P. Pentland, Face recognition using eigenfaces, in: Computer Vision and Pattern Recognition, 1991. Proceedings CVPR'91., IEEE Computer Society Conference on, IEEE, 1991, pp. 586–591.
- [37] P. Viola, M.J. Jones, Robust real-time face detection, Int. J. Comput. Vision 57 (2) (2004) 137–154.
- [38] L. Wang, Y. Qiao, X. Tang, Action recognition with trajectory-pooled deepconvolutional descriptors, arXiv preprint arXiv:1505.04868(2015).
- [39] Z. Wang, A.C. Bovik, H.R. Sheikh, E.P. Simoncelli, Image quality assessment: from error visibility to structural similarity, IEEE Trans. Image Process. 13 (4) (2004) 600–612.
- [40] L. Wolf, T. Hassner, Y. Taigman, The one-shot similarity kernel, in: 12th International Conference on Computer Vision, IEEE, 2009, pp. 897–902.
- [41] D.H. Wolpert, Stacked generalization, Neural Netw. 5 (2) (1992) 241-259.
- [42] J. Wright, A.Y. Yang, A. Ganesh, S.S. Sastry, Y. Ma, Robust face recognition via sparse representation, Pattern Anal. Mach. Intell. IEEE Trans. 31 (2) (2009) 210–227.
- [43] T. Wu, R. Chellappa, Age invariant face verification with relative craniofacial growth model, in: Computer Vision-ECCV 2012, 2012, pp. 58–71.
- [44] D. Yi, Z. Lei, S. Liao, S.Z. Li, Learning face representation from scratch, arXiv preprint arXiv:1411.7923(2014).
- [45] J. Yosinski, J. Clune, Y. Bengio, H. Lipson, How transferable are features in deep neural networks? in: Advances in Neural Information Processing Systems, 2014, pp. 3320–3328.
- [46] S. Zha, F. Luisier, W. Andrews, N. Srivastava, R. Salakhutdinov, Exploiting image-trained cnn architectures for unconstrained video classification, arXiv preprint arXiv:1503.04144(2015).
- [47] A. Krizhevsky, I. Sutskever, G.E. Hinton, Imagenet classification with deep convolutional neural networks, Advances in neural information processing systems, 2012, pp. 1097–1105.