

Age Estimation Based on Face Images and Pre-trained Convolutional Neural Networks

Abhinav Anand, Ruggero Donida Labati, Angelo Genovese,

Enrique Muñoz, Vincenzo Piuri, Fabio Scotti

Department of Computer Science, Università degli Studi di Milano, Italy.

firstname.lastname@unimi.it

Abstract—Age estimation based on face images plays an important role in a wide range of scenarios, including security and defense applications, border control, human-machine interaction in ambient intelligence applications, and recognition based on soft biometric information. Recent methods based on deep learning have shown promising performance in this field. Most of these methods use deep networks specifically designed and trained to cope with this problem. There are also some studies that focus on applying deep networks pre-trained for face recognition, which perform a fine-tuning to achieve accurate results. Differently, in this paper, we propose a preliminary study on increasing the performance of pre-trained deep networks by applying post-processing strategies. The main advantage with respect to fine-tuning strategies consists of the simplicity and low computational cost of the post-processing step. To the best of our knowledge, this paper is the first study on age estimation that proposes the use of post-processing strategies for features extracted using pre-trained deep networks. Our method exploits a set of pre-trained Convolutional Neural Networks (CNNs) to extract features from the input face image. The method then performs a feature level fusion, reduces the dimensionality of the feature space, and estimates the age of the individual by using a Feed-Forward Neural Network (FFNN). We evaluated the performance of our method on a public dataset (Adience Benchmark of Unfiltered Faces for Gender and Age Classification) and on a dataset of non-ideal samples affected by controlled rotations, which we collected in our laboratory. Our age estimation method obtained better or comparable results with respect to state-of-the-art techniques and achieved satisfactory performance in non-ideal conditions. Results also showed that CNNs trained on general datasets can obtain satisfactory accuracy for different types of validation images, also without applying fine-tuning methods.

I. INTRODUCTION

Face biometrics is the most natural method for human recognition because of its high level of user acceptance due to the low level of required user cooperation [1], [2]. Face images can also be used to infer a wide set of soft biometric characteristics, such as the emotional state, ethnicity, gender, and age. Among this set of characteristics, the automatic age estimation can be particularly important in different scenarios [3], such as security and defense applications [4], surveillance [5], health-care systems [6], [7], entertainment [8], automated

border controls [9], [10], human-machine interactions in ambient intelligence applications [11], and recognition based on soft biometric traits [12].

Recent works in the literature showed that deep learning techniques can achieve promising accuracy in age estimation [13], [14]. The studies described in [15]–[17] analyzed the possibility of estimating the age by using deep networks previously trained for face recognition. However, all of these methods apply fine-tuning strategies to achieve accurate results on heterogeneous image datasets.

This paper presents a preliminary study to increase the accuracy of pre-trained deep networks for age estimation without applying fine-tuning approaches. To the best of our knowledge, this paper presents the first study on age estimation that uses previously trained Convolutional Neural Networks (CNNs) without needing any training or fine-tuning of the deep neural networks, thus considering CNNs trained for other applications as generic feature extractors. This approach has the advantage of simplifying the post-processing task with respect to fine-tuning techniques.

We propose the use of heterogeneous networks trained using non-ideal samples to extract robust features from non-ideal face images. Our method performs a feature level fusion of the data computed by a set of CNNs, reduces the dimensionality of the obtained feature set, and estimates the age by using a Feed-Forward Neural Network (FFNN). To achieve robustness to non-ideal conditions, we tune the dimensionality reduction method using public face datasets of poor-quality face images acquired in uncontrolled conditions.

The contribution of the paper is three-fold. First, we present a novel age estimation method that robustly extracts features from non-ideal face images. Second, we propose a feature level strategy to fuse data extracted by heterogeneous CNNs. Third, we prove that CNNs previously trained for different applications can be used as general feature extractors for estimating soft-biometric information from non-ideal face samples, also without applying fine-tuning techniques.

We evaluated the accuracy of our method for a dataset of 4535 non-ideal face images collected in our laboratory and compared its performance with other methods in the literature using the Adience Benchmark of Unfiltered Faces for Gender and Age Classification [14]. Results proved the effectiveness of our method for non-ideal face images, thus showing that

This work was supported in part by: the EC within the 7FP under grant agreement 312797 (ABC4EU); the EC within the H2020 program under grant agreement 644597 (ESCUDO-CLOUD); and the Italian Ministry of Research within PRIN 2015 project COSMOS (201548C5NT).




Fig. 1. Schema of the proposed approach for age estimation.

CNNs trained on general datasets can obtain satisfactory accuracy for different types of validation images even without applying fine-tuning techniques. Our method achieved better or in-line accuracy with respect to state-of-the-art techniques based on feature sets tuned for the Adience Benchmark of Unfiltered Faces for Gender and Age Classification.

This paper is structured as follows: Section II presents a brief literature review on age estimation methods from face images, Section III describes the proposed method, Section IV presents the performed experiments, and Section V concludes the work.

II. LITERATURE REVIEW

Facial aging is a widely studied problem. In the literature, there are techniques for age estimation [18], simulation of the aging process [19], and biometric recognition designed to cope with images acquired at different ages [20].

The studies in the literature on age estimation methods can be divided into machine learning approaches based on feature sets designed to extract discriminative age characteristics and methods that directly infer knowledge from the samples by using deep neural networks.

In the first category, there are studies focusing on general texture features, such as Local Binary Patterns [21], [22] and Gabor features [23]. Other studies focus on computing novel features specifically designed for age estimation [24], [25]. The approach presented in [24] uses bio-inspired features based on pyramids of Gabor filters. The method described in [26] uses the Active Appearance Model to extract the regions of age local features. Recent techniques also propose ad-hoc learning methods able to exploit information from the ordinal relationship of the aging labels [27], [28].

In the second category, pattern recognition techniques based on deep learning have shown promising results in many real world applications [29], [30]. The research community has also successfully designed and trained deep neural networks for age estimation [13], [14], [31], [32]. Most of these techniques are based on deep CNNs [31], [32]. Other methods are based on different approaches, such as the group-aware deep feature learning [13] or drop-out Support Vector Machines [14]. These studies achieved high accuracy for heterogeneous datasets acquired in different application scenarios. However, the main drawback consists of the relevant amount of time needed to train deep neural networks in different application scenarios.

Some studies [15]–[17] investigated the possibility of estimating the age by using a deep CNN previously trained for

face recognition applications. However, all of these methods perform a fine-tuning of the CNN, thus requiring a considerable amount of images and time for being tuned and applied in new scenarios.

Differently from other approaches based on deep learning techniques, our method uses previously trained CNNs as general feature extractors, but without training or fine-tuning deep neural networks. Our method is thus easier to be applied in novel conditions because it is based on a simpler and faster learning process.

III. PROPOSED METHOD

The proposed method estimates the age from a single face image by using multiple pre-trained deep networks that permit to robustly estimate features from non-ideal face images. Our method can be divided into three main steps: *i*) feature extraction using pre-trained CNNs, *ii*) dimensionality reduction, and *iii*) age estimation using FFNNs. Fig. 1 shows the schema of the proposed age estimation method.

A. Feature extraction using CNNs

We use one or more pre-trained CNNs to extract features from the input face images for age estimation. In this paper, we use VGG-Face CNN [33] and AlexNet CNN [34], although these two deep networks can be easily substituted.

VGG-Face CNN consists of 15 layers, trained on 2.6M facial images from 2622 individuals to perform face recognition. Each CNN block contains a linear operator followed by one or more non-linear layers, such as a rectification layer or max pooling. The first 12 such blocks are convolutional layers. The last 3 blocks are fully connected. We used the second fully connected layer for feature extraction, obtaining 4096 dimensional feature sets.

AlexNet CNN consists of 8 layers, trained on 1.2M samples of ImageNet ILSVRC challenge dataset [35]. As in VGG-Face CNN, each CNN block contains a linear operator followed by one or more non-linear layers. The first 5 blocks are convolutional layers, while the last 3 are fully connected. We used the second fully connected layer for feature extraction, obtaining 4096 dimensional feature sets.

B. Dimensionality reduction

First, we create a feature set of 8192 features by using the data extracted using VGG-Face and AlexNet CNNs. Then, we apply a dimensionality reduction technique. We analyzed the results of three strategies, namely: Statistical Dependency algorithm (SD) [36], Mutual Information (MI)

[36], and Principal Component Analysis (PCA) [37]. To obtain an age estimation method that is easily applicable in different application scenarios, we choose to adopt an external image dataset to train the parameters of the evaluated dimensionality reduction methods only once. In particular, to increase the robustness of our method to rotations and non-ideal samples, we train the dimensionality reduction strategies using a public dataset of poor-quality face images acquired in uncontrolled conditions (WIKI Dataset [31]).

- **Statistical Dependency (SD):** this method aims to measure the statistical dependency of the features to its class labels. The feature set is first quantized in such a way that each bin contains roughly the same number of samples from the whole dataset. The statistical dependency between the feature values x and the age y , modeled as a class instead of a continuous value, is computed as:

$$SD = \sum_{x \in X} \sum_{y \in Y} p(x, y) \frac{p(x, y)}{p(x)p(y)} . \quad (1)$$

We calculate the SD for each feature and select the features having the highest values as final feature set.

- **Mutual Information (MI):** this method quantifies the mutual dependency between the two random variables. This method takes into account that the statistical dependency between the features and the age labels may be affected by the highly informative level of quantization [36]. To reduce this effect, we compute the MI between the features and the age labels as:

$$MI = \sum_{x \in X} \sum_{y \in Y} p(x, y) \log \left(\frac{p(x, y)}{p(x)p(y)} \right) . \quad (2)$$

- **Principal Component Analysis (PCA):** this method computes the Eigenvectors from the covariance matrix. The covariance matrix evaluates the amount by which neighboring features are related to each other. The Eigenvectors associated with the largest Eigenvalue are the ones that reflect the largest variance in the image.

C. Age estimation using FFNNs

The last step uses the features computed in the dimensionality reduction step to train a FFNN. We use a FFNN for regression to estimate a floating point number representing the age and similarly we apply a FFNN for classifying groups of ages. In particular, we use a single linear node for the output layer of the neural network and the scaled conjugate gradient backpropagation method [38] for training.

IV. EXPERIMENTAL RESULTS

In this section, we discuss the used datasets, the configuration of the experiments, the results obtained by the proposed method, and the required computational time. In particular, we analyze three test scenarios to evaluate our method. First, we evaluate the performance of the studied dimensionality reduction and fusion techniques. Second, we

present the performance of our method for a wide set of non-idealities. Third, we compare the performance of our method with that of state-of-the-art techniques on a challenging public dataset.

A. Database description

We used public biometric datasets as well as sets of samples acquired in our laboratory by simulating the acquisitions performed in less-constrained conditions. We tested our method using the following datasets:

- **WIKI Dataset:** we used this database [31] to estimate the parameters of the feature selection methods and to train the FFNNs. It consists of 62,359 images of popular celebrities from Wikipedia. Most of the images in the database display variations of poses and illumination conditions.
- **Aml-Face Dataset:** we created this dataset to evaluate our method in non-ideal conditions. It consists of face images captured at various distances (25 cm, 50 cm, 100 cm, and 150 cm) using a webcam. To simulate a less-constrained acquisition scenario, we acquired faces at various degrees of rotations such as 0°, 22°, 45°, 75°, and 90°. Furthermore, we included two additional activity scenarios: in the first scenario, the images are acquired when the user is talking on the smartphone; in the second scenario, faces are acquired with strong expression differences. The dataset is composed of 4535 face images of 16 individuals. We captured these images using a Microsoft LifeCam HD-3000. Fig. 2 shows some examples of collected face images. To process these samples using CNNs, we cropped the face region using the face detector proposed by Viola and Jones [39].
- **Adience Benchmark Dataset:** to compare the classification accuracy of our method with state of the art techniques, we used the dataset Adience Benchmark of Unfiltered Faces for Gender and Age Classification [14]. This dataset consists of 26,000 face images from 2284 individuals. The dataset is categorized into eight age groups: {[0, 2], [4, 6], [8, 13], [15, 20], [25, 32], [38, 43], [48, 53], [60, -]}. In our study, we assigned integer classes from 1 to 8 to the age groups sorted in increasing order. As described in [14], to process these samples using CNNs, we cropped the face region using the face detector proposed by Viola and Jones [39] and aligned using the frontalization method described in [40].

B. Configuration of the experiments

We tested different numbers of nodes with tan-sigmoidal transfer functions for the hidden layer of the FFNN regression models (in the range of [1-100]). We imposed a maximum of 2000 training epochs. We obtained the best results using 15 nodes in the hidden layer. We evaluated the age estimation error in terms of MAE (Mean Absolute Error), defined as the average of the absolute difference between the estimated age and the actual age. When testing the age classification accuracy, we evaluated the exact classification error.




Fig. 2. Examples of cropped faces in Aml-Face Dataset simulating less-constrained and non-cooperation scenarios including rotations: (a, h) frontal, (b, i) 22°, (c, j) 45°, (d, k) 75°, (e, l) 90°; and activities: (f, m) using smartphone, and (g, n) expression changes.

C. Evaluation of dimensionality reduction and fusion methods

We evaluated the accuracy of the different dimensionality reduction methods. To analyze the capability of our method to be applied in different scenarios without performing any training, we trained the dimensionality reduction and regression models using the WIKI Dataset, then we tested the obtained models on the Aml-Face dataset. In particular, we checked the performance of our method using the features obtained from VGG-Face CNN, AlexNet CNN, and with different dimensionality reduction strategies. We also compared the performance of the proposed feature level fusion with that of the single CNNs and with a method performing a decision level fusion (computed as the mean of the ages estimated using the CNNs singularly). Table I summarizes the obtained results.

Results show that, for all the tested configurations, the dimensionality reduction improved the estimation of MAE with respect to the full feature set. In our tests, VGG-Face CNN performed better than AlexNet CNN, probably due to the fact that VGG-Face CNN has been designed for face recognition and AlexNet CNN has been designed for analyzing general images. The best performing feature set is the feature level fusion of VGG-Face CNN and AlexNet CNN using 30 features obtained after applying PCA. The PCA is the dimensionality reduction method that achieved the best accuracy, probably due to its intrinsic capability of reducing noise by discarding the less significant Eigenvectors. Another important observation is that fusion strategies always increased the performances of single CNNs. In particular, the feature level fusion obtained better results than the single CNNs and the decision level fusion.

D. Analysis of the performance for simulated less-constrained acquisition scenarios

We evaluated separately the accuracy of the best configuration of our method for each non-ideality of the Aml-Face Dataset. Table II summarizes the achieved results. This table

shows that the higher performance decrease corresponds to the scenario in which the users were speaking using a smartphone and to the scenario in which the users were appositely performing strong changes in their expression. In all these cases, MAE decreased by 0.46 years with respect to acquisitions performed with frontal face and neutral expression. We believe that this performance decrease can be considered as satisfactory for age estimation methods working in unconstrained scenarios. It is also interesting to note that MAE decreased by 0.46 years for acquisition performed with head rotations of 90° with respect to frontal acquisitions. This result is particularly promising since most of the methods in the literature are designed to work with frontal acquisitions and have not been evaluated with this kind of strong face rotations. The obtained performance suggests that groups of CNNs trained using big and heterogeneous datasets can extract discriminative features robust to non-ideal application conditions.

E. Comparison with recent methods in the literature

To evaluate the performance of our method in general scenarios and compare our results with the state of the art methods, we used the Adience Benchmark Dataset. We adopted the 5-fold cross-validation procedure suggested in [14] and evaluated the performance in terms of exact classification accuracy across all age groups.

Table III compares the accuracy of the best configuration of our method with that of state-of-the-art techniques for the Adience Benchmark Dataset. All the reported results have been computed using the same 5-fold cross-validation procedure. For this scenario, the best configuration consisted of the feature level fusion of 500 values obtained by applying PCA to the feature sets extracted using VGG-Face CNN and AlexNet CNN, respectively.

Table III shows that our method obtained better or comparable classification accuracy with respect to state-of-the-art methods based on specifically designed features or on deep

TABLE I
COMPARISON OF DIMENSIONALITY REDUCTION AND FUSION METHODS IN TERMS OF MAE (IN YEARS)

Features	Test database											
	CNNi + PCA + FFNN				CNNi + MI + FFNN				CNNi + SD + FFNN			
	V	A	Mean	V+A	V	A	Mean	V+A	V	A	Mean	V+A
10	4.38	4.72	3.44	3.51	7.03	4.41	5.12	—	6.96	4.41	5.07	—
30	4.05	5.17	3.87	3.30	7.93	4.28	5.14	9.86	7.05	4.21	4.82	9.39
50	4.17	6.58	4.33	3.40	5.46	5.33	4.14	7.39	4.80	5.31	4.13	7.39
100	4.10	5.09	3.64	3.45	4.70	5.23	3.94	—	4.68	5.51	4.08	—
200	4.52	5.74	3.88	3.85	4.40	5.56	5.12	—	4.13	4.89	3.49	—
All	V (4096 features)				A (4096 features)				V+A (8192 features)			
	5.4				5.86				—			

We trained the dimensionality reduction and the regression methods using the WIKI Dataset, while we tested the performance of our method on the AmI-Face dataset. V = VGG-Face CNN, A = AlexNet CNN, Mean = decision level fusion of V and A, V+A = feature level fusion of V and A, — = the FFNN did not converge to a suitable model.

TABLE II
PERFORMANCE OF THE PROPOSED METHOD IN LESS-CONSTRAINED CONDITIONS IN TERMS OF MAE (IN YEARS)

Frontal	22°	45°	75°	90°	Smartp.	Expr.	DB
3.12	3.03	3.00	3.23	3.48	3.58	3.58	3.30

We trained the feature reduction and the regression methods using the WIKI dataset, while we tested the performance of our method on the AmI-Face dataset. Smartp. = smartphone, Expr. = expression changes, and DB = complete dataset.

networks trained for the considered dataset. The table also shows that our method achieved better accuracy with respect to most of the techniques based on the fine-tuning of pre-trained deep networks, thus proving the feasibility of using pre-trained CNNs as generic feature extractors for age estimation.

F. Analysis of the computational time

This section analyzes the time required by the steps of our method. We executed the tests using a laptop PC with 2.2 GHz Intel (R) i7 (R) CPU, RAM 8 GB and NVIDIA GeForce GT 640M GPU. The operating system was Windows 7 professional 64 bit. All methods were implemented using Matlab. The CNNs used the GPU parallelization implemented by the library MatConvNet [44]. The total time required to estimate the age from each face image was 0.59 s. VGG-Face CNN required 0.51 s, while VGG-Face CNN required 0.07 s, and the dimensionality reduction and classification steps required a total time of 0.01 s.

We think that the computational cost of our method is acceptable for online applications. Furthermore, we expect that compiled languages, such as C/C++, can reduce the processing time with respect to the considered implementation.

Since VGG-Face CNN required most of the computational time, less complex CNNs could be considered for application scenarios characterized by particularly strict constraints in terms of computational time.

V. CONCLUSION AND FUTURE WORK

In this paper, we presented a novel age estimation method designed to increase the performance of pre-trained deep networks by applying post-processing strategies. Our method uses pre-trained sets of CNNs to extract features from the face images. It then performs a dimensionality reduction of the feature set, a feature level fusion, and estimates the age value using feedforward neural networks.

We evaluated our method by training it on a public dataset and testing it on images acquired in a less-constrained scenario. The achieved results show that the our method achieved satisfactory performance for non-ideal images acquired in unconstrained scenarios. We also compared the accuracy of our age estimation method with that of state-of-the-art techniques by using a challenging public dataset. The obtained results show that our method achieved better or comparable results with respect to the state of the art. Results also demonstrated that CNNs trained on general datasets can obtain satisfactory accuracy for different types of validation images. Furthermore, results proved that pre-trained deep networks can be considered as general feature extractors for age estimation, also without applying computationally expensive fine-tuning techniques.

As future work, we should test the suitability of the proposed method for estimating other soft biometric information, such as gender or pose. Moreover, we should evaluate other feature reduction strategies widely used by face recognition systems.

REFERENCES

- [1] A. Jain, P. Flynn, and A. Ross, "Handbook of biometrics," *Springer Science & Business Media*, 2007.
- [2] R. Donida Labati, A. Genovese, E. Muñoz, V. Piuri, F. Scotti, and G. Sforza, "Biometric recognition in automated border control: A survey," *ACM Computing Surveys*, vol. 49, no. 2, pp. 24:1–24:39, November 2016.
- [3] A. Dantcheva, P. Elia, and A. Ross, "What else does your biometric data reveal? a survey on soft biometrics," *IEEE Trans. on Information Forensics and Security*, vol. 11, no. 3, pp. 441–467, September 2016.

TABLE III
RESULTS IN LITERATURE FOR AGE CLASSIFICATION ON ADIENCE BENCHMARK DATASET USING CLASSIFICATION ACCURACY

Method	Year	Description	Classification accuracy (%)
[14]	2014	Dropout SVM + LBP	45.10
[32]	2015	Deep CNN	50.70
[41]	2017	Deep CNN	61.30
[42]	2017	Deep identification-verification features (DeepID2)	51.10
[43]	2017	Feedforward attention mechanism	61.80
[15]	2016	Pre-trained CNN + fine-tuning	57.90
[16]	2016	Pre-trained CNN + fine-tuning	64.00
[17]	2016	Pre-trained CNN + fine-tuning	52.88
Proposed method	2017	Multiple pre-trained CNNs	58.49

Tests performed using Adience benchmark dataset with 5-fold cross validation method. We used the parameters of PCA estimated using WIKI Dataset.

- [4] Q. Xiao, "A biometric authentication approach for high security ad-hoc networks," in *Proc. of the Fifth Annual IEEE SMC Information Assurance Workshop (IAW)*, 2004, pp. 250–256.
- [5] P. Tome, J. Fierrez, R. Vera-Rodriguez, and M. S. Nixon, "Soft biometrics and their application in person recognition at a distance," *IEEE Trans. on Information Forensics and Security*, vol. 9, no. 3, pp. 464–475, March 2014.
- [6] Q. Zhang, D. Zhou, and X. Zeng, "HeartID: A multiresolution Convolutional Neural Network for ECG-based biometric human identification in smart health applications," *IEEE Access*, vol. 5, pp. 11 805–11 816, 2017.
- [7] R. Donida Labati, V. Piuri, R. Sassi, G. Sforza, and F. Scotti, "Adaptive ECG biometric recognition: a study on re-enrollment methods for QRS signals," in *Proc. of the IEEE Workshop on Computational Intelligence in Biometrics and Identity Management (CIBIM)*, 2014, pp. 30–37.
- [8] L. Geng, K. Zhang, X. Wei, and X. Feng, "Soft biometrics in online social networks: A case study on Twitter user gender recognition," in *Proc. of the 2017 IEEE Winter Applications of Computer Vision Workshops (WACVW)*, 2017, pp. 1–8.
- [9] A. Anand, R. Donida Labati, A. Genovese, E. Muñoz, V. Piuri, F. Scotti, and G. Sforza, "Enhancing the performance of multimodal Automated Border Control systems," in *Proc. of the 2016 Int. Conf. of the Biometrics Special Interest Group (BIOSIG)*, 2016, pp. 1–5.
- [10] A. Anand, R. Donida Labati, A. Genovese, E. Muñoz, V. Piuri, F. Scotti, and G. Sforza, "Enhancing fingerprint biometrics in Automated Border Control with adaptive cohorts," in *Proc. of the 2016 IEEE Symp. on Computational Intelligence for Security and Defense Applications (CISDA)*, 2016, pp. 1–8.
- [11] A. Anand, R. Donida Labati, M. Hanmandlu, V. Piuri, and F. Scotti, "Text-independent speaker recognition for Ambient Intelligence applications by using information set features," in *Proc. of the 2017 IEEE Int. Conf. on Computational Intelligence and Virtual Environments for Measurement Systems and Applications (CIVEMSA)*, 2017, pp. 30–35.
- [12] R. Donida Labati, A. Genovese, V. Piuri, and F. Scotti, "Weight estimation from frame sequences using computational intelligence techniques," in *Proc. of the 2012 IEEE Int. Conf. on Computational Intelligence for Measurement Systems and Applications (CIMSAS)*, 2012, pp. 29–34.
- [13] H. Liu, J. Lu, J. Feng, and J. Zhou, "Group-aware deep feature learning for facial age estimation," *Pattern Recognition*, vol. 66, no. Supplement C, pp. 82–94, June 2017.
- [14] E. Eiding, R. Enbar, and T. Hassner, "Age and gender estimation of unfiltered faces," *IEEE Trans. on Information Forensics and Security*, vol. 9, no. 12, pp. 2170–2179, December 2014.
- [15] G. Ozbulak, Y. Aytar, and H. K. Ekenel, "How transferable are CNN-based features for age and gender classification?" in *Proc. of the 2016 IEEE Int. Conf. of the Biometrics Special Interest Group (BIOSIG)*, 2016, pp. 1–6.
- [16] R. Rothe, R. Timofte, and L. Van Gool, "Deep expectation of real and apparent age from a single image without facial landmarks," *Int. Journal of Computer Vision*, pp. 1–14, August 2016.
- [17] J. C. Chen, A. Kumar, R. Ranjan, V. M. Patel, A. Alavi, and R. Chellappa, "A cascaded Convolutional Neural Network for age estimation of unconstrained faces," in *Proc. of the 8th IEEE Conf. on Biometrics Theory, Applications and Systems (BTAS)*, 2016, pp. 1–8.
- [18] T. Dhiman and K. Mistree, "Feature extraction for facial age estimation: A survey," in *Proc. of the 2016 Int. Conf. on Wireless Communications, Signal Processing and Networking (WiSPNET)*, 2016, pp. 2243–2248.
- [19] A. Lanitis, C. J. Taylor, and T. F. Cootes, "Toward automatic simulation of aging effects on face images," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 24, no. 4, pp. 442–455, April 2002.
- [20] S. Sahni and S. Saxena, "SURVEY: Techniques for aging problems in face recognition," *MIT Int. Journal of Computer Science and Information Technology*, vol. 4, no. 2, pp. 82–88, August 2014.
- [21] J. Ylioinas, A. Hadid, X. Hong, and M. Pietikäinen, "Age estimation using Local Binary Pattern kernel density estimate," in *Proc. of the 17th International Conference on Image Analysis and Processing (ICIAP)*, 2013, pp. 141–150.
- [22] A. Gunay and V. Nalbant, "Automatic age classification with LBP," in *Proc. of the 23rd IEEE Int. Symp. on Computer and Information Sciences (ISCIS)*, 2008, pp. 1–4.
- [23] F. Gao and H. Ai, "Face age classification on consumer images with Gabor feature and fuzzy LDA method," in *Proc. of the Int. Conf. on Biometrics (ICB)*, 2009, pp. 132–141.
- [24] G. Guo, G. M. Y. Fu, and T. S. Huang, "Human age estimation using bio-inspired features," in *Proc. of the 2009 IEEE Int. Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2009, pp. 112–119.
- [25] X. Geng, Z. Zhou, and K. Smith-Miles, "Automatic age estimation based on facial aging patterns," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 29, no. 12, pp. 2234–2240, December 2007.
- [26] J. D. Tria and C. L. Huang, "Age estimation using AAM and local facial features," in *Proc. of the 5th Int. Conf. on Intelligent Information Hiding and Multimedia Signal Processing (IIH-MSP)*, 2009, pp. 885–888.
- [27] X. Geng, C. Yin, and Z. Zhou, "Facial age estimation by learning from label distributions," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 35, no. 10, pp. 2401–2412, October 2013.
- [28] S. Feng, C. Lang, J. Feng, T. Wang, and J. Luo, "Human facial age estimation by cost-sensitive label ranking and trace norm regularization," *IEEE Trans. on Multimedia*, vol. 19, no. 1, pp. 136–148, January 2017.
- [29] R. Donida Labati, A. Genovese, E. Muñoz, V. Piuri, and F. Scotti, "A novel pore extraction method for heterogeneous fingerprint images using Convolutional Neural Networks," *Pattern Recognition Letters*, 2017.
- [30] A. Toshev and C. Szegedy, "DeepPose: Human pose estimation via Deep Neural Networks," in *Proc. of 2014 IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2014, pp. 1653–1660.
- [31] R. Rothe, R. Timofte, and L. Van Gool, "Dex: Deep expectation of apparent age from a single image," in *Proc. of the 2015 IEEE Int. Conf. on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2015, pp. 10–15.
- [32] G. Levi and T. Hassner, "Age and gender classification using Convolutional Neural Networks," in *Proc. of the 2015 IEEE Int. Conf. on*

- Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2015, pp. 34–42.
- [33] O. M. Parkhi, A. Vedaldi, and A. Zisserman, “Deep face recognition,” in *Proc. of the British Machine Vision Conference (BMVC)*, 2015.
 - [34] A. Krizhevsky, I. Sutskever, and G. E. Hinton, “ImageNet classification with Deep Convolutional Neural Networks,” *Advances in Neural Information Processing Systems* 25, pp. 1097–1105, 2012.
 - [35] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. M. Z. Huang, A. Karpathy, A. Khosla, and M. Bernstein, “ImageNet large scale visual recognition challenge,” *Int. Journal of Computer Vision*, vol. 115, no. 3, pp. 211–252, December 2015.
 - [36] J. Pohjalainen, O. Räsänen, and S. Kadioglu, “Feature selection methods and their combinations in high-dimensional classification of speaker likability, intelligibility and personality traits,” *Computer Speech & Language*, vol. 29, no. 1, pp. 145–171, January 2015.
 - [37] A. Malhi and R. X. Gao, “PCA-based feature selection scheme for machine defect classification,” *IEEE Trans. on Instrument and Measurements*, vol. 53, no. 6, pp. 1517–1525, December 2004.
 - [38] M. F. Moller, “A scaled conjugate gradient algorithm for fast supervised learning,” *Neural Networks*, vol. 6, no. 4, pp. 525–533, 1993.
 - [39] P. Viola and M. Jones, “Rapid object detection using a boosted cascade of simple features,” in *Proc. of the 2001 IEEE Computer Society Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2001, pp. 511–518.
 - [40] T. Hassner, S. Harel, E. Paz, and R. Enbar, “Effective face frontalization in unconstrained images,” in *Proc. of the 2015 IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2015, pp. 4295–4304.
 - [41] A. Dehghan, E. G. Ortiz, G. Shu, and S. Z. Masood, “DAGER: Deep Age, Gender and Emotion Recognition using Convolutional Neural Network,” *CoRR*, vol. abs/1702.04280, 2017.
 - [42] J. Huang, B. Li, J. Zhu, and J. Chen, “Age classification with Deep Learning face representation,” *Multimedia Tools and Applications*, vol. 76, no. 19, pp. 20 231–20 247, October 2017.
 - [43] P. Rodríguez, G. Cucurull, J. M. Gonfaus, F. X. Roca, and J. González, “Age and gender recognition in the wild with deep attention,” *Pattern Recognition*, vol. 72, no. Supplement C, pp. 563–571, December 2017.
 - [44] A. Vedaldi and K. Lenc, “Matconvnet – convolutional neural networks for matlab,” in *Proceeding of the ACM Int. Conf. on Multimedia*, 2015.