Combination of Texture and Geometric Features for Age Estimation in Face Images

Marcos Vinicius Mussel Cirne and Helio Pedrini
Institute of Computing, University of Campinas, Campinas, SP, 13083-852, Brazil

Keywords: Age Estimation, Image Analysis, Texture, Geometric Descriptor.

Abstract: Automatic age estimation from facial images has recently received an increasing interest due to a variety of applications, such as surveillance, human-computer interaction, forensics, and recommendation systems. Despite such advances, age estimation remains an open problem due to several challenges associated with the aging process. In this work, we develop and analyze an automatic age estimation method from face images based on a combination of textural and geometric features. Experiments are conducted on the Adience dataset (Adience Benchmark, 2017; Eidinger et al., 2014), a large known benchmark used to evaluate both age and gender classification approaches.

1 INTRODUCTION

Biometric systems commonly employ a number of distinctive measurable human characteristics to recognize individuals, such as face, fingerprint, palm print, deoxyribonucleic acid (DNA) and iris (Paulo Carlos et al., 2015; Pinto et al., 2015; Silva Pinto et al., 2012; Menotti et al., 2015; Menotti et al., 2015; Silva et al., 2015; Assis Angeloni and Pedrini, 2016). Soft biometrics (Jain et al., 2004) refer to metrics related to physical or behavioral human characteristics, such as gender, age, hair color, height and weight, which are complementary to primary biometric identifiers. The combination of primary and soft biometric characteristics can significantly improve the performance of person recognition in surveillance systems.

The problem of age estimation from face images (Choi et al., 2011; Huerta et al., 2014; Lanitis et al., 2002; Liu et al., 2015; Ren and Li, 2014; Thukral et al., 2012; Geng et al., 2013) is very challenging due to the inherently complex nature of the aging process, high variability within a same age interval, personal characteristics of each individual, as well as difficulties in collecting large datasets derived from chronological images from the same individuals.

Despite its large applicability in several knowledge domains, there is still relatively little research on age estimation compared to other facial analysis topics, such as face and iris recognition. Examples of applications of age estimation from facial images include forensics, surveillance, human-computer interaction, and recommendation systems.

In this work, we propose and evaluate a novel age estimation approach from facial images based on a combination of textural and geometric features.

Experiments are conducted on the Adience dataset (Adience Benchmark, 2017; Eidinger et al., 2014), which is a well known benchmark used to evaluate age estimation approaches.

The remainder of this paper is organized as follows. A literature review about age estimation approaches is shown in Section 2. The proposed age estimation method is detailed in Section 3. Experimental results are presented and discussed in Section 4. Finally, some final remarks and directions for future work are included in Section 5.

2 BACKGROUND

Due to a variety of applications, there has been an increasing interest of the scientific community in investigating the automatic age estimation from facial images.

Most of the age estimation approaches available in the literature are based on feature extraction and learning algorithms. A pioneer work on age estimation from facial images was proposed by Kwon and Lobo (Kwon and da Vitoria Lobo, 1994), which was based on an anthropomorphic model to differentiate

Cirne, M. and Pedrini, H.
Combination of Texture and Geometric Features for Age Estimation in Face Images.
DOI: 10.5220/0006625503950401
In *Proceedings of the 13th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications (VISIGRAPP 2018) - Volume 4: VISAPP*, pages 395-401
ISBN: 978-989-758-290-5

three age groups: babies, young adults and senior adults. In the same work, they also proposed a wrinkle detection method based on snakelets.

Lanitis et al. (Lanitis et al., 2002) employed active appearance models (AAM) for age estimation by defining an ageing function. Chang et al. (Chang et al., 2011) also used AAM to estimate ordinal hyperplanes and ranks them according to age intervals. Geng et al. (Geng et al., 2013) explored an ageing pattern subspace model to extract and process feature vector for age estimation.

Fu and Huang (Fu and Huang, 2008) developed a manifold embedding approach to the age estimation problem, whose purpose is to find a low-dimensional representation in the embedded subspace and capture geometric structure and data distribution. Wu et al. (Wu et al., 2012) explored a Grassmann manifold to model the facial shapes and considered the age estimation as regression and classification problems on this representation.

Appearance models have also been explored by several authors for age estimation purpose. Gao and Ai (Gao and Ai, 2009) used a texture descriptor based on Gabor filter to estimate age. Guo and Guowang (Guo and Mu, 2014), Guo et al. (Guo et al., 2009) and Weng et al. (Weng et al., 2013) employed biologically inspired features for age estimation from face images.

Hayashi et al. (Hayashi et al., 2002) developed a method for age estimation based on wrinkle texture and color information extracted from facial images. Shape and size of the facial parts were used to predict the age.

Iga et al. (Iga et al., 2003) described extraction functions from facial candidate regions using color information and parts of the face. Age was estimated based on SVM classifiers and a voting scheme from the extracted features.

Suo et al. (Suo et al., 2010) proposed a compositional and dynamic model for age estimation, where a hierarchical and-or graph was used to represent faces in each age group. A Markov process was employed to parse the graph representation.

Kawano et al. (Kawano et al., 2005) described a four-directional feature based on multiple parts of the face, such as nose, lip, jaw, eyes. Linear discriminant analysis was employed to recognize these image regions.

Luu et al. (Luu et al., 2011) developed an appearance model based on contourlet transforms to locate facial landmarks. Local and holistic texture information was explored for facial age estimation.

For further details on age estimation models and algorithms, the reader can refer to surveys by Dhimar and Mistree (Dhimar and Mistree, 2016) and Fu et al. (Fu et al., 2010).

3 PROPOSED METHOD

The proposed approach to the age classification problem is divided into two major steps: pre-processing stage and cross-validation stage. Figure 1 shows an overview of these stages.

In the pre-processing stage, a dataset containing several images of people of different ages (or intervals of ages) is analyzed. After that, the feature extraction process is started. For each image of the dataset, several image descriptors are extracted, providing a "descriptor database" that allows the use of different combinations of descriptors.

For this work, both image and geometric features were taken into account for estimating ages from images. Concerning image features, two descriptors were extracted from the images: the first one is the Histogram of Oriented Gradients (HOG) (Dalal and Triggs, 2005), a very popular and robust descriptor used for detection and recognition of objects and faces (Déniz et al., 2011; Felzenszwalb et al., 2010; Suard et al., 2006; Zhu et al., 2006).

First, each image was downsampled to a size of 64×64 pixels in order to produce a more compact descriptor. To calculate the HOG descriptor, the images were initially split into 8×8 cells. Then, for the block normalization process, a block size of 16×16 with a stride of 8×8 was used, thus producing 7 horizontal and 7 vertical positions, resulting 49 positions. By defining a histogram size of 9 bins, since each block encompasses 4 cells and each cell provides a single histogram, each step of the block normalization process creates a feature vector of 36 dimensions. Hence, the final HOG descriptor has a size of $36 \times 49 = 1764$ dimensions.

The second image descriptor used in our method was the CENTRIST (CENsus TRansform hIStogram) (Wu and Rehg, 2011), an evolution of the LBP-texture descriptor (Ojala et al., 2002) that encodes the structural properties of an image at the same time it has a high level of robustness to illumination variations. Here, for each pixel x_c of an grayscale image, its intensity is compared against its 8 neighbors x_p, where $p = 0, 1, ...7$, producing a bit string $B_c = b_0 b_1 ... b_7$. For each comparison, if $x_c \geq x_p$, then we have $b_p = 1$. Otherwise, $b_p = 0$. The final bit string is then converted to a base-10 value in the range $[0, 255]$.

Equation 1 shows the pattern used in this process, along with an example. Once this procedure is

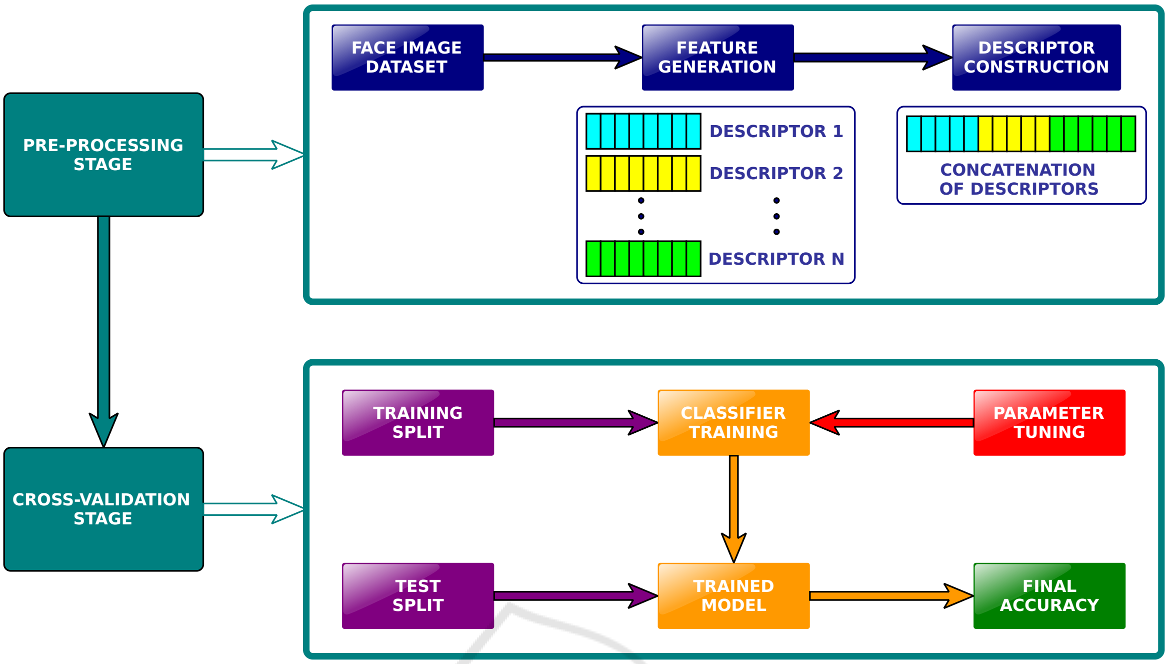

Figure 1: General scheme of our approach to age classification using combinations of image descriptors and machine learning techniques.

done for all pixels, a histogram of 256 bins is computed from the computed values, which forms the final CENTRIST descriptor.

$$\begin{bmatrix} x_0 & x_1 & x_2 \\ x_3 & x_c & x_4 \\ x_5 & x_6 & x_7 \end{bmatrix} \Rightarrow \begin{bmatrix} 40 & 160 & 120 \\ 180 & 100 & 80 \\ 100 & 20 & 80 \end{bmatrix} \quad (1)$$

$$\Rightarrow \quad B_c = 01110100_2 = 116_{10}$$

Besides image features, a geometric descriptor was also included in the set of descriptors. Given a pre-trained face shape model and an image from the age dataset, a face detection algorithm is run to find both the region that covers the face and, if a face is found, the facial landmarks that correspond to important facial features (such as eyes, nose, mouth, eyebrows) according to the shape model.

In this work, the algorithm and the model provided by the Dlib toolkit (Dlib Toolkit, 2017) were used in this process. The algorithm for face detection is an implementation based on the work by Kazemi and Sullivan (Kazemi and Sullivan, 2014), which uses an ensemble of regression trees to estimate facial landmarks from a sparse subset of pixel intensities in a fast and efficient way.

The shape model used in this process, along with the localization and indexing of the facial landmarks, is shown in Figure 2. Based on this model, the face detection algorithm extracts the exact localization of the pixels that correspond to each of the 68 points of the model. Once the set of points is obtained, a geometric descriptor is constructed by taking the Euclidean distances between all pairs of points, which gives a total of

$$\binom{68}{2} = 2278 \text{ distances.}$$

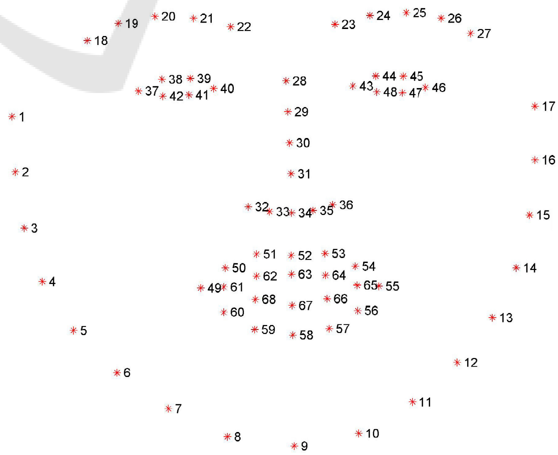

Figure 2: Facial landmarks of the face shape model used in Dlib's face detection algorithm (Dlib Toolkit, 2017).

After extracting several image descriptors, the descriptor construction process is started by creating different concatenations of descriptors that will represent each image from a dataset. Since the features of each descriptor are from different natures and sizes, the values of the final descriptor are rescaled so that they have zero mean and unit variance. In this work, 7 different combinations of descriptors were used, which will be detailed later on Section 4.

Once the pre-processing stage is finished, the dataset is split into training and test sets using the k-fold cross-validation. In this procedure, the entire dataset is split into k parts such that one of these parts is left apart for the tests and the $k-1$ remaining ones are used for training a classifier. Later, the test set is used to evaluate the accuracy of the trained model and the entire process is repeated for the other folds such that each fold is used once as the test set, which gives a total of k iterations. Finally, the mean accuracy and the standard error of the trained model are calculated from the accuracy rates obtained in each iteration of the cross-validation process.

For this stage of the proposed method, the multi-class version of the Support Vector Machines (SVM) classifier (Cortes and Vapnik, 1995; Hearst et al., 1998) with an RBF kernel was used. Other possibilities of classifiers (Alpaydin, 2014; Bishop, 2006; Murphy, 2012) were also considered, such as K-Nearest Neighbors, Logistic Regression and Random Forests, but the SVM produced the best results among all those options.

4 EXPERIMENTAL RESULTS

The proposed method was implemented using Scikit-Learn (Scikit-Learn Machine Learning in Python, 2017), an open source Python library that contains several tools for machine learning, image processing, data mining and data analysis. From the set of descriptors detailed in Section 3 (HOG, CENTRIST and the geometric descriptor), all possible combinations among them were evaluated.

The tests were conducted on the Adience dataset (Adience Benchmark, 2017; Eidinger et al., 2014), created by the Open University of Israel (OUI) to facilitate the study of both age and gender classification problems. The collection contains over 26,580 images of 2,284 subjects from 8 different age groups, from newborns to old-aged people. All photos were taken with several variations in appearance, posing, lighting, background and facial expressions. Besides the original images, a special version containing cropped and aligned face images is also available. In addition, the dataset provides information for a 5-fold cross-validation procedure, listing the images that make part of each split. Figure 3 shows some examples from the dataset.

Figure 3: Examples of images from the Adience dataset (Eidinger et al., 2014).

For this work, only frontal images from the aligned version of the dataset were used, since the face detection algorithm does not perform very well with rotated faces (i.e., face images that are out of the $\pm 5°$ range of yaw angle), making the computation of the geometric descriptor difficult for these cases. After generating all descriptors for the remaining images, the dataset was reduced to a total of 11,437 images. A complete list of the amount of images in each fold for each age interval can be seen in Table 1.

The evaluation process is the same adopted by Eidinger et al. (Eidinger et al., 2014). Here, two different accuracies are calculated: the exact accuracy, which computes the groups that were correctly predicted, and the 1-off accuracy, which also considers errors of one age group as correct predictions (e.g.: for a face image whose correct class is “15-20”, “8-12” and “25-32” predictions are also regarded as correct). Then, the mean accuracies of all 5 folds of the database and their respective standard errors are calculated. Table 2 shows the results for all combinations of descriptors, where the boldface values represent the best results obtained by our approach.

From the table, it can be seen that the three descriptors combined produced the best results among all possibilities for both exact and 1-off accuracy rates. Analyzing each descriptor separately, the HOG descriptor provided the highest accuracy rates. Moreover, when concatenating a different descriptor to an existing one, both accuracies are increased. In other words, the descriptor formed by the combination of the three descriptors used in this work performs better than all combinations of two descriptors, which

Table 1: Amounts of aligned frontal face images from the Adience dataset (Adience Benchmark, 2017; Eidinger et al., 2014) for each fold and age interval.

Fold #	**Age Intervals**								**Total**
	0-2	4-6	8-12	15-20	25-32	38-43	48-53	60+	
Fold 1	642	354	153	105	1062	381	158	88	2943
Fold 2	155	326	414	355	415	288	84	92	2129
Fold 3	608	256	335	157	481	195	70	113	2215
Fold 4	108	195	389	300	570	312	65	77	2016
Fold 5	195	381	212	142	573	305	152	174	2134
Total	1708	1512	1503	1059	3101	1481	529	544	11437

Table 2: List of mean exact and 1-off accuracies for different combinations of image descriptors for the Adience dataset (Adience Benchmark, 2017; Eidinger et al., 2014).

Descriptor(s)	**Exact Acc. (%)**	**1-Off Acc. (%)**
CENTRIST	30.38 ± 6.71	54.84 ± 5.19
Geometric	41.10 ± 4.76	76.55 ± 2.09
HOG	43.69 ± 5.70	77.03 ± 2.09
CENTRIST + Geometric	43.22 ± 5.67	78.51 ± 2.90
CENTRIST + HOG	45.03 ± 7.00	78.08 ± 2.66
Geometric + HOG	45.99 ± 6.07	81.07 ± 2.26
CENTRIST + Geometric + HOG	$\mathbf{46.70 \pm 6.56}$	$\mathbf{81.80 \pm 2.23}$
LBP + FBLBP (Eidinger et al., 2014)	44.5 ± 2.3	80.7 ± 1.1
LBP + FBLBP + Dropout-SVM (Eidinger et al., 2014)	45.1 ± 2.6	79.5 ± 1.4

in turn perform better than using only one descriptor, except for the exact accuracy in the HOG versus CENTRIST + Geometric case.

Considering the maximum accuracies obtained, these results are slightly better than the ones achieved by Eidinger et al. (Eidinger et al., 2014), who obtained accuracies of 44.5 ± 2.3 and 80.7 ± 1.1 for the exact and 1-off cases, respectively, only using variations of the LBP descriptor and linear SVM. However, they managed to improve the exact accuracy to 45.1 ± 2.6 by using a variation of SVM called dropout-SVM, inspired by the concept of dropout from deep neural networks (Krizhevsky et al., 2012). Therefore, our approach has great space for improvements, not only in the choice of image descriptors, but also in the machine learning process as a whole.

5 CONCLUSIONS AND FUTURE WORK

This work presented an approach to age estimation from face images using image descriptors that encompass both texture and geometric features. The Adience dataset was used as the case study and trained with an SVM classifier with an RBF kernel.

Tests were conducted with several combinations of descriptors, also observing how the addition of a specific descriptor affects the mean accuracies. Results showed that our approach is comparable to the LBP-feature based approach (Eidinger et al., 2014), even though a different classifier was used to achieve the best results.

Future directions for this work include: tests with new sets of descriptors (including a refinement of the geometric descriptor), use of different age datasets and benchmarks and parameter optimization for the machine learning procedure, application of deep learning techniques to improve the quality of the results, along with an evaluation of other types and variations of classifiers.

ACKNOWLEDGMENTS

The authors are thankful to São Paulo Research Foundation (grants FAPESP #2017/12646-3 and

#2014/12236-1) and National Council for Scientific and Technological Development (grant CNPq #305169/2015-7) for their financial support.

REFERENCES

Adience Benchmark (2017). Unfiltered Faces for Gender and Age Classification. http://www.openu.ac.il/home/hassner/Adience/data.html.

Alpaydin, E. (2014). *Introduction to Machine Learning*. MIT Press.

Assis Angeloni, M. and Pedrini, H. (2016). Part-based Representation and Classification for Face Recognition. In *IEEE International Conference on Systems, Man, and Cybernetics*, pages 2900–2905, Budapest, Hungary. IEEE.

Bishop, C. M. (2006). *Pattern Recognition and Machine Learning*. Springer.

Chang, K.-Y., Chen, C.-S., and Hung, Y.-P. (2011). Ordinal Hyperplanes Ranker with Cost Sensitivities for Age Estimation. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 585–592. IEEE.

Choi, S. E., Lee, Y. J., Lee, S. J., Park, K. R., and Kim, J. (2011). Age Estimation using a Hierarchical Classifier based on Global and Local Facial Features. *Pattern Recognition*, 44(6):1262–1281.

Cortes, C. and Vapnik, V. (1995). Support-Vector Networks. *Machine Learning*, 20(3):273–297.

Dalal, N. and Triggs, B. (2005). Histograms of Oriented Gradients for Human Detection. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, volume 1, pages 886–893.

Déniz, O., Bueno, G., Salido, J., and la Torre, F. D. (2011). Face recognition using Histograms of Oriented Gradients. *Pattern Recognition Letters*, 32(12):1598–1603.

Dhimar, T. and Mistree, K. (2016). Feature Extraction for Facial Age Estimation: A Survey. In *International Conference on Wireless Communications, Signal Processing and Networking*, pages 2243–2248.

Dlib Toolkit (2017). http://dlib.net.

Eidinger, E., Enbar, R., and Hassner, T. (2014). Age and Gender Estimation of Unfiltered Faces. *IEEE Transactions on Information Forensics and Security*, 9(12):2170–2179.

Felzenszwalb, P. F., Girshick, R. B., McAllester, D., and Ramanan, D. (2010). Object Detection with Discriminatively Trained Part-Based Models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32(9):1627–1645.

Fu, Y., Guo, G., and Huang, T. S. (2010). Age Synthesis and Estimation via Faces: A Survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32(11):1955–1976.

Fu, Y. and Huang, T. S. (2008). Human Age Estimation with Regression on Discriminative Aging Manifold. *IEEE Transactions on Multimedia*, 10(4):578–584.

Gao, F. and Ai, H. (2009). Face Age Classification on Consumer Images with Gabor Feature and Fuzzy LDA Method. In *International Conference on Biometrics*, pages 132–141. Springer.

Geng, X., Yin, C., and Zhou, Z.-H. (2013). Facial Age Estimation by Learning from Label Distributions. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(10):2401–2412.

Guo, G. and Mu, G. (2014). A Framework for Joint Estimation of Age, Gender and Ethnicity on a Large Database. *Image and Vision Computing*, 32(10):761–770.

Guo, G., Mu, G., Fu, Y., and Huang, T. S. (2009). Human Age Estimation using Bio-inspired Features. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 112–119. IEEE.

Hayashi, J.-i., Yasumoto, M., Ito, H., and Koshimizu, H. (2002). Age and Gender Estimation based on Wrinkle Texture and Color of Facial Images. In *16th International Conference on Pattern Recognition*, volume 1, pages 405–408. IEEE.

Hearst, M. A., Dumais, S. T., Osuna, E., Platt, J., and Scholkopf, B. (1998). Support Vector Machines. *IEEE Intelligent Systems and their Applications*, 13(4):18–28.

Huerta, I., Fernández, C., and Prati, A. (2014). Facial Age Estimation through the Fusion of Texture and Local Appearance Descriptors. In *European Conference on Computer Vision*, pages 667–681. Springer International Publishing.

Iga, R., Izumi, K., Hayashi, H., Fukano, G., and Ohtani, T. (2003). A Gender and Age Estimation System from Face Images. In *SICE Annual Conference*, volume 1, pages 756–761. IEEE.

Jain, A. K., Dass, S. C., and Nandakumar, K. (2004). Soft Biometric Traits for Personal Recognition Systems. In *Biometric Authentication*, pages 731–738. Springer.

Kawano, T., Kato, K., and Yamamoto, K. (2005). An Analysis of the Gender and Age Differentiation using Facial Parts. In *IEEE International Conference on Systems, Man and Cybernetics*, volume 4, pages 3432–3436. IEEE.

Kazemi, V. and Sullivan, J. (2014). One Millisecond Face Alignment with an Ensemble of Regression Trees. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1867–1874. IEEE Computer Society.

Krizhevsky, A., Sutskever, I., and Hinton, G. E. (2012). ImageNet Classification with Deep Convolutional Neural Networks. In Pereira, F., Burges, C. J. C., Bottou, L., and Weinberger, K. Q., editors, *Advances in Neural Information Processing Systems*, volume 25, pages 1097–1105. Curran Associates, Inc.

Kwon, Y. H. and da Vitoria Lobo, N. (1994). Age Classification from Facial Images. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 762–767. IEEE.

Lanitis, A., Taylor, C. J., and Cootes, T. F. (2002). Toward Automatic Simulation of Aging Effects on Face

Images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(4):442–455.

Liu, K.-H., Yan, S., and Kuo, C.-C. J. (2015). Age Estimation via Grouping and Decision Fusion. *IEEE Transactions on Information Forensics and Security*, 10(11):2408–2423.

Luu, K., Seshadri, K., Savvides, M., Bui, T. D., and Suen, C. Y. (2011). Contourlet Appearance Model for Facial Age Estimation. In *International Joint Conference on Biometrics*, pages 1–8. IEEE.

Menotti, D., Chiachia, G., Pinto, A., Schwartz, W. R., Pedrini, H., Falcão, A. X., and Rocha, A. (2015). Deep Representations for Iris, Face, and Fingerprint Spoofing Detection. *IEEE Transactions on Information Forensics and Security*, 10(4):864–879.

Murphy, K. P. (2012). *Machine Learning: A Probabilistic Perspective*. MIT Press.

Ojala, T., Pietikäinen, M., and Mäenpää, T. (2002). Multiresolution Gray-Scale and Rotation Invariant Texture Classification with Local Binary Patterns. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(7):971–987.

Paulo Carlos, G., Pedrini, H., and Schwartz, W. R. (2015). Classification Schemes based on Partial Least Squares for Face Identification. *Journal of Visual Communication and Image Representation*, 32:170–179.

Pinto, A., Pedrini, H., Schwartz, W. R., and Rocha, A. (2015). Face Spoofing Detection through Visual Codebooks of Spectral Temporal Cubes. *IEEE Transactions on Image Processing*, 24(12):4726–4740.

Ren, H. and Li, Z.-N. (2014). Age Estimation based on Complexity-aware Features. In *Asian Conference on Computer Vision*, pages 115–128. Springer.

Scikit-Learn Machine Learning in Python (2017). http://scikit-learn.org/stable/.

Silva, P., Luz, E., Baeta, R., Pedrini, H., Falcão, A. X., and Menotti, D. (2015). An Approach to Iris Contact Lens Detection based on Deep Image Representations. In *28th SIBGRAPI Conference on Graphics, Patterns and Images*, pages 157–164. IEEE.

Silva Pinto, A., Pedrini, H., Schwartz, W., and Rocha, A. (2012). Video-based Face Spoofing Detection through Visual Rhythm Analysis. In *25th SIBGRAPI Conference on Graphics, Patterns and Images*, pages 221–228. IEEE.

Suard, F., Rakotomamonjy, A., Bensrhair, A., and Broggi, A. (2006). Pedestrian Detection using Infrared images and Histograms of Oriented Gradients. In *IEEE Intelligent Vehicles Symposium*, pages 206–212.

Suo, J., Zhu, S.-C., Shan, S., and Chen, X. (2010). A Compositional and Dynamic Model for Face Aging. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32(3):385–401.

Thukral, P., Mitra, K., and Chellappa, R. (2012). A Hierarchical Approach for Human Age Estimation. In *IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 1529–1532. IEEE.

Weng, R., Lu, J., Yang, G., and Tan, Y.-P. (2013). Multi-feature Ordinal Ranking for Facial Age Estimation. In *IEEE International Conference and Workshops on Automatic Face and Gesture Recognition*, pages 1–6. IEEE.

Wu, J. and Rehg, J. M. (2011). CENTRIST: A Visual Descriptor for Scene Categorization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(8):1489–1501.

Wu, T., Turaga, P., and Chellappa, R. (2012). Age Estimation and Face Verification across Aging using Landmarks. *IEEE Transactions on Information Forensics and Security*, 7(6):1780–1788.

Zhu, Q., Yeh, M.-C., Cheng, K.-T., and Avidan, S. (2006). Fast Human Detection Using a Cascade of Histograms of Oriented Gradients. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, volume 2, pages 1491–1498.