HAL
archives-ouvertes.fr

Cross-database evaluation of normalized raw pixels for gender recognition under unconstrained settings

Taner Danisman, Ioan Marius Bilasco, Chaabane Djeraba

▶ **To cite this version:**

Taner Danisman, Ioan Marius Bilasco, Chaabane Djeraba. Cross-database evaluation of normalized raw pixels for gender recognition under unconstrained settings. 22nd International Conference on Pattern Recognition (ICPR 2014), Aug 2014, Stockholm, Sweden. pp.3144-3149, 10.1109/ICPR.2014.542 . hal-00973505

HAL Id: hal-00973505
https://hal.archives-ouvertes.fr/hal-00973505

Submitted on 28 Sep 2018

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Cross-database evaluation of normalized raw pixels for gender recognition under unconstrained settings

Taner Danisman, Ioan Marius Bilasco, Chabane Djeraba
Lille 1 University, LIFL
IRCICA, Parc Scientifique de la Haute Borne
50 Avenue Halley 59655 Villeneuve d'Ascq - France
{taner.danisman,marius.bilasco,chabane.djeraba}@lifl.fr

***Abstract*—This paper presents cross-database evaluations of automatic appearance-based gender recognition methodology using normalized raw pixels and SVM classifier under unconstrained settings. Proposed method uses both histogram specification and feature space normalization on automatically aligned faces to achieve reliable recognition rate for real scenarios. Using a web based unconstrained training database, we applied local window search to increase generalization ability of the proposed method. Our contribution is two-fold. First we showed that aligned and normalized raw pixel intensities are providing the best performance in case of unconstrained cross-database tests than feature-based studies on unaligned faces. Second, we showed that histogram specification provides better normalization than that of histogram equalization for automatically aligned faces in large databases for gender recognition. Variety of cross-database experiments performed on uncontrolled Image of Groups (88.16%), Genki-4K (91.07%) and LFW databases (91.87%) showed that proposed method provides superior generalization ability than that of the state-of-the-art methods.**

I. Introduction

Gender recognition is a challenging task in computer vision. It is a necessary component of demographic studies performing gender, age and ethnicity analysis. Since it is mainly a pattern recognition problem, advancements in this task can also be applicable to other pattern recognition problems. For this reason gender recognition is getting more interest from computer vision domain. Due to its intra-class variability, gender recognition is a special case of a two-class classification task. This variation mainly comes from different age, occlusions, ethnicity, make-up, illumination and data acquisition conditions.

Previous studies are mainly focused on controlled single-database experiments with classical five-fold cross validation technique which, due to acquisition protocols of the controlled databases, may lead to biased results. In general, as indicated by Bekios-Calfa et al. [1], the performance of intra-database experiments is optimistic because of the fact that images in these databases usually have a common demography and acquisition protocol. According to Boon et al. [2], average accuracy in gender recognition studies performed on controlled Feret database [3] is $96\% \pm 2.5$ especially when a cross validation scheme is applied. However, this high performance decrease drastically when a cross-database scheme used $83\% \pm 6.7$. Several researchers provide cross-database evaluation of their methods performed in unconstrained Labeled Faces in the Wild [4] and Image of Groups [5] databases. In these uncontrolled databases, the average cross-database accuracy is even lower $79.05\% \pm 5.6$ see Table I) showing the lack of inconsistency of the generalization ability of the models across different databases. Therefore future research must focus more on improving the generalization ability of the methodologies under unconstrained settings than controlled settings. Detailed comparisons with the state-of-the-art methods are provided in discussion section IV.

Initial studies in the domain considered appearance based features like raw pixels [6], [7]. More recent studies focus on feature based methods [8]–[11], histogram oriented gradients (HOG) and Gabor filters for unconstrained gender recognition [9]. LBP operator [12] and its variants are also widely used in feature based methods. A recent survey explaining board range of methodologies for vision based gender recognition is presented in [2]. According to the literature survey, SVM classifier with RBF kernel is the most common classifier used in gender recognition studies because of its generalization ability. Other methods include Adaboost, nearest neighbor classifier and neural networks.

Makinen et al. [13] studied the performance of automatically aligned faces with non-aligned and manually aligned faces for gender recognition in Feret database to measure robustness of these preprocessing methods along with the state-of-the-art methods. They concluded that the automatic alignment does not increase the gender classification rates and that pixel-based input gives better classification than LBP-based features. However, their experiments performed on a small subset (60 male, 47 female) of controlled Feret database. We believe that controlled databases and their subsets are not representative for reliable recognitions in gender recognition problem. Since the training set is just a small part of the whole gender recognition problem (7 billion people), classifiers learned from a very specific demographic subset do not perform well on unseen faces. Therefore, a successful gender recognition methodology must consider unconstrained settings when selecting the preprocessing methodology. Our experiments on relatively large-scale and unconstrained LFW and Image of Groups databases showed that specific age subsets of the faces tend to be classified better than others. In turn it affects the overall classification performance of the classifier. Therefore the evaluation and interpretation of gender recognition methodologies depend on overall demographic diversity of the training and test database.

In this study, we compared the cross-database performance

of automatically aligned faces using feature space normalization with SVM+RBF classifier to the previously published results on these unconstrained databases. Main contribution of this paper is twofold. First, we showed that aligned and normalized raw pixel intensities provide the best performance in case of unconstrained cross-database tests than LBP based features even in larger subsets of the databases. Second, we showed that histogram specification provides better normalization than that of histogram equalization for automatically aligned faces in large databases for gender recognition.

Section 2 explains the methodology and databases used in the experiment. Section 3 presents experimental setup and results. Section 4 provides comparison and discussion of the obtained results and previous studies in the domain. Final section summarizes and concludes the study with future directions.

II. Methodology

First, we collected a new gender database using web images by querying search engines. In order to get wide range of ethnicity, we used common "male face" and "female face" query words in different languages including Arabic, Chinese, English, French, German, Italian, Spanish, Russian and Turkish. Manually eliminating the false results, we obtained a new unconstrained database including low resolution 3630 female and 1692 male images. Image size is varies between 60×80 and 300×400. We combined this new database with manually annotated Genki-4K database [14] to obtain a large-scale training database WebDB having 5640 females and 3686 male samples.

A. Preprocessing

The first step towards unconstrained gender recognition is preprocessing available information in proper way. Geometric and photometric normalization methods are of crucial importance for vision based facial analysis. Normalized information provides better approximation of noisy data to the mathematical models used in computer vision.

We start with the detection of face using well-known "frontal_alt2" haar-like features model [15] available in OpenCV [16]. After that, eye detection is performed to correct in-plane rotation of the face according to the vertical position of left and right pupil. We used the neural network-based eye detector [17] available in the Stacked Trimmed Active Shape Model (STASM) [18] library to locate the positions of the pupils.

1) Face Alignment: Updated face rectangle is defined using the following equations where F_x, F_y, F_w and F_h represents the x, y, $width$ and $height$ of the updated face. Eye_{Left_x} and Eye_{Left_y} are the x and y positions of the left eye with respect to upper left origin of the image. IPD (Inter-Pupillary Distance) is Euclidean distance between the eye centers.

$$F_x = Eye_{Left_x} - IPD/4.0 \quad (1)$$

$$F_y = Eye_{Left_y} - IPD \quad (2)$$

$$F_w = IPD \times 1.5 \quad (3)$$

$$F_h = IPD \times 2.5 \quad (4)$$

Note that initial location of the OpenCV face detection results are updated according to the IPD distance. Scalar values 4.0, 1.5 and 2.5 are selected according to experimental observations. Extracted face is then resized to 20×24 image. Finally, histogram specification is applied to overcome illumination differences. Figure 1 shows the initial and cropped face region after the use of Eq. (1), (2), (3) and (4).



Fig. 1. Geometric normalization and alignment of the face

2) Histogram Specification: Histogram specification and histogram equalization are fundamental image enhancement techniques used in image processing. Histogram equalization assigns equal number of pixels to all gray levels. However, this method does not consider common facial appearance. Histogram specification is a generalization of histogram equalization where the image is normalized with respect to a desired probability density function (pdf). Since we know an average human face, we can apply the histogram extracted from the average face to all normalized images. Figure 2 shows the effect of histogram specification for a given face image using the histogram of the average face.

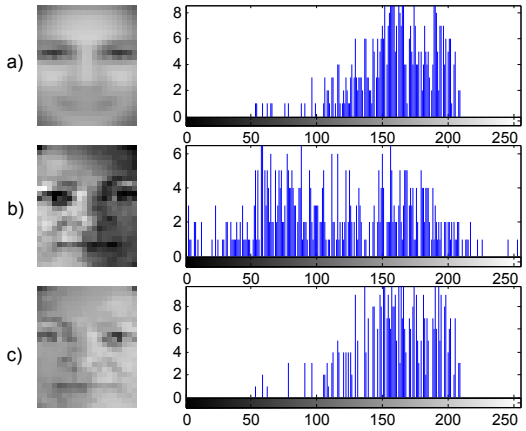


Fig. 2. a) Average image obtained from web database and its histogram. b) Example test image and corresponding histogram. c) Result of histogram specification on b) using the histogram of a)

As seen on Figure 2 (c), estimated new histogram is more close to the histogram of the average face. This feature provides better correction of the image histogram in case of different illumination conditions as seen on Figure 3.

3) Feature Space Normalization: Arnulf et al. [19] showed that preprocessing step can be equivalent to studying the kernel functions of SVMs. In addition they concluded that feature space normalization outperforms input space normalization in

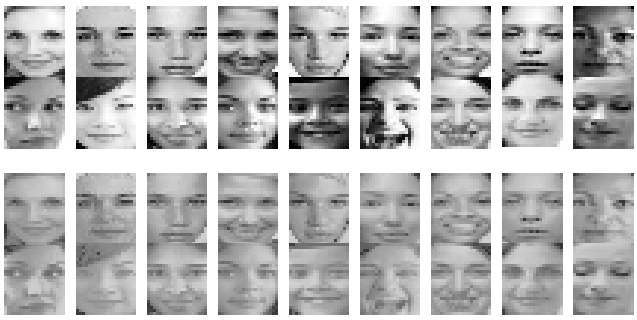


Fig. 3. Normalized face samples of size 20×24 (first two row) and result of histogram specification (last two row)

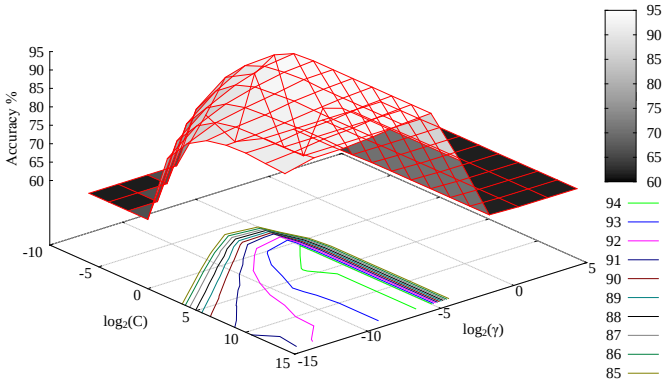


Fig. 4. SVM+RBF parameter space and corresponding five-fold cross-validation accuracy for WebDB. Selected parameters $C = 16$ ($\log_2 C = 4$) and $\gamma = 0.0078125$ ($\log_2 \gamma = -7$)

case of SVMs. Therefore we employed feature space normalization on normalized faces obtained in the previous step. Since we have $20 \times 24 = 480$ features for a single image, feature space normalization performed on each of the 480 features. Considering all training database samples, we found the minimum and maximum value in each feature and the new value of the feature is computed between -1 and +1. However, when the training size increases, possibility of having the smallest and largest pixel value (0-255) also increases. Another say, feature space normalization become equivalent to non-normalized methods when individual scaling factors per feature are similar.

III. Experimental Setup

The experiments have been designed to test the generalization ability of the SVM+RBF classifier trained on normalized raw pixels using unconstrained WebDB, LFW, Genki-4K and Groups databases. In order to apply a consistent protocol, we used the same classification parameters in all experiments. There is a notable unbalance between female and male classes in LFW database (2978 female, 10258 male). Therefore, we focused mainly on Groups database. Due to the difficulties of comparing different cross-database evaluations, we provide all details of the experimental setup.

A. Parameter Selection

After the preprocessing step, normalized faces of size 20×24 are used for both training the SVM with RBF kernel and selection of optimal cost (C) and gamma (γ) parameters. We used five-fold cross-validation method on WebDB database for the parameter selection by $easy$ tool present in LibSVM [20] where a grid search is applied. Since the combination of large γ and large C leads overfitting, we selected $C = 16$ and $\gamma = 0.0078125$ as the optimum values which are similar to the results obtained by Makinen and Raisamo [21]. This setting provides 94.27% accuracy using five-fold cross-validation on our web database WebDB. Figure 4 shows the parameter space and corresponding accuracies obtained from five-fold cross validation.

B. Databases

We selected the LFW, Genki-4K and Groups database which are the most unconstrained databases and therefore giving the lowest accuracies in cross-database tests. WebDB is used for parameter selection and training the model as described in section III-A.

1) Image of Groups (Groups): Groups database [5] includes 5080 images having 28231 faces labeled with the age and gender categories. It involves wide range of illumination, ethnicity, ages, facial expressions, in-plane and out-of-plane poses. Manually labeled eye positions are provided for all faces. However, we automatically detect the faces and eyes using the methods provided in section II-A. We obtained a total of 19835 faces (10303 female, 9532 male).

2) Labeled Faces in the Wild (LFW): LFW database [4] contains 13236 labeled images from 5749 individuals mainly actors, politicians and sport players. We automatically select detected faces where eye detection is successful (11106) and then manually group them into male (8539) and female (2567) categories.

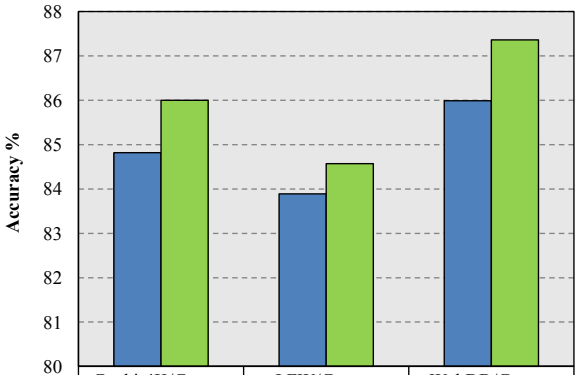
3) Genki-4K Subset: Genki-4K database mainly used in facial expression studies and it contains 4000 face images labeled as either smiling or non-smiling. It involves wide range of subjects, facial appearance, illumination, geographical locations, imaging conditions, and camera model. However, it does not include gender labels. Therefore we manually label the images as female and male classes for our experiments. After the normalization step we obtained 1539 females and 1506 males.

IV. Results and Discussion

We performed series of experiments on WebDB, LFW, Genki-4K and Groups databases comparing the effect of using histogram specification and histogram equalization. We also studied the relationship between age and gender by cross-database tests. Finally, we compared our cross-database results with the state-of-the-art methods.

A. Histogram Specification Experiments

We compared the results obtained from histogram equalization and histogram specification on Groups database. For training we used Genki-4K, LFW and WebDB databases. Figure 5 shows effect of using histogram equalization and histogram specification on Groups database where age of the subject is greater than 12.



Train/Test	Genki-4K/Groups	LFW/Groups	WebDB/Groups
Hist. Equalization	84.82	83.89	85.99
Hist. Specification	86.00	84.57	87.36

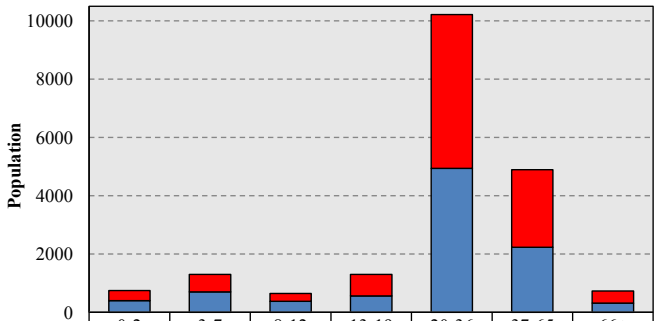
Fig. 5. Comparison of histogram equalization and histogram specification on Groups database (age > 12) using cross-database experiments.

According to Figure 5 we obtain better classification results when using histogram specification method in training and testing. We believe that main reason of this improvement (1%) is better correction of illumination by considering the pdf of the average face.

B. Gender and Age Experiments

Previously Guo et al. [22] showed that gender recognition accuracies can be 10% higher on adult faces than on young or senior faces. Their study on a large database also showed that gender recognition is still affected significantly when using either basic raw pixels or more advanced features like HOG, BIF or LBP descriptors.

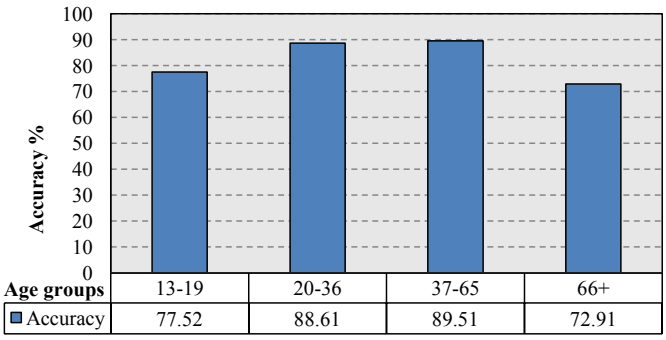
Groups database provides age labels in 0-2, 3-7, 8-12, 13-19, 20-36, 37-65 and 65+ age categories. Since our training sets Genki-4K, LFW and WebDB does not contain children faces, for this experiment we present our results by considering age categories greater than 12. However, we also provide global result including all age categories which is 82.09% as shown in Table I. Figure 6 shows age histogram of Groups database. It is clear that adult face categories are dominant than young and senior categories.



Age Groups	0-2	3-7	8-12	13-19	20-36	37-65	66+
Female	346	604	265	738	5278	2659	413
Male	405	701	381	561	4935	2230	318

Fig. 6. Age histogram of Groups database.

Figure 7 shows the accuracy of WebDB/Groups experiment while considering different age groups in Groups database. As we expected, the best results obtained from adult category 20-36 and 37-65 with 88.61% and 89.51% accuracy. We obtained better results on teenager category 13-19 than senior category 66+ which are 77.52% and 72.91% respectively. Main reason to get better results on adult category might be the demographic similarity of WebDB which consist of only adult faces. Since adult faces are also dominant in LFW and Genki-4K databases, we obtained similar results from Genki-4K/Groups and LFW/Groups experiments.



Age groups	13-19	20-36	37-65	66+
Accuracy	77.52	88.61	89.51	72.91

Fig. 7. Accuracy for different age groups. Training with WebDB and testing in Groups database.

C. Cross-database Experiments

We present state-of-the-art gender recognition studies and our results using cross-database evaluations in Table I. Highlighted entries present the best accuracy obtained so far in corresponding test database. We obtained 88.16% accuracy on Groups for adult people having age ≥ 20, 87.36% for age> 12. These results are superior to that of the state-of-the-art results as shown in Table I. Similarly we obtained 91.87% accuracy on WebDB/LFW and 91.07% on Groups/Genki-4K tests.

As shown in Table II, it is notable that all cross-database results are balanced in the range 85.00% and 92.05% which is an indicator of a good generalization level for the proposed method. In addition we did not restrict our test sets for a minimum IPD distance. For this reason we obtain larger test sets than those used in state-of-the-art studies. As shown in Table I, we used 19834 test samples from Groups database without any age and IPD distance limitation, 17132 test samples for age > 12 and 15833 test samples for age ≥ 20. Study by Bekios-Calfa et al. [23] considers unaligned faces for gender recognition. When they consider different pose as labeled powerset LP [24], obtained 79.53% and 78.33% with and without pose factor respectively. However, these results are lower than 91.62% accuracy obtained from our Groups/LFW experiment.

Considering Table II, we strictly separate face of the same identity appearing both in training and test sets. Otherwise, this may lead the classifier learning the identity instead of the gender. Since Genki-4K is a subset of WebDB as described in section II, we did not perform any experiment between Genki-4K and WebDB.

TABLE I. SUMMARY OF CROSS-DATABASE EXPERIMENTS ON GROUPS, LFW AND GENKI-4K DATABASES.

Author	Train/Test	Test size	Method	Acc.%
Dago-Casas [11]	LFW/Groups[a]	14760	LBP+PCA+LDA	81.02
Dago-Casas [11]	LFW/Groups[a]	14760	Pixels+PCA+SVM	72.09
Balmaseda [10]	Morph[25]/Groups	5632	LBP+SVM+Linear	76.74
Balmaseda [10]	Morph/LFW[b]	1149	LBP+SVM+Linear	75.10
Dago-Casas [11]	Groups/LFW	13088	LBP+PCA+SVM	89.77
Bekios-Calfa [23]	Groups/LFW	13233	PCA+LDA+KNN	79.11
Bekios-Calfa [23]	Groups[c]/LFW	13233	PCA+LDA+KNN	79.53
Our method	WebDB/Groups	**19834**	Pixels+SVM+RBF	**82.09**
Our method	WebDB/Groups[d]	17132	Pixels+SVM+RBF	87.36
Our method	WebDB/Groups[e]	15833	Pixels+SVM+RBF	**88.16**
Our method	Genki-4K/Groups[e]	15833	Pixels+SVM+RBF	86.78
Our method	LFW/Groups[d]	17132	Pixels+SVM+RBF	84.57
Our method	LFW/Groups[e]	15833	Pixels+SVM+RBF	85.00
Our method	LFW/Genki-4K	3045	Pixels+SVM+RBF	87.62
Our method	Groups/Genki-4K	3045	Pixels+SVM+RBF	**91.07**
Our method	WebDB/LFW	11106	Pixels+SVM+RBF	**91.87**

[a] inter-ocular distance ≥ 20 pixels
[b] one image per identity
[c] training step without children faces
[d] ages > 12
[e] ages ≥ 20

TABLE II. SUMMARY OF CROSS-DATABASE EXPERIMENTS ON ALL DATABASES (ACCURACY %).

Train \ Test	Genki-4K	Groups	LFW	WebDB
Genki-4K	×	86.78	88.18	×
Groups	91.07	×	91.62	88.70
LFW	87.62	85.00	×	92.05
WebDB	×	88.16	91.87	×

V. CONCLUSION

In this study we present unconstrained cross-database experiments on publicly available gender databases LFW, Groups and Genki-4K using histogram specification based normalization and SVM+RBF kernel. We show that histogram specification provides better normalization than histogram equalization in case of unconstrained databases. Using aligned and normalized 20×24 raw pixels, we obtained better generalization and consistent results than existing state-of-the-art methods performed on cross-database tests. In addition we studied the effect of age on gender recognition and showed that adult gender recognition rate is higher than that of teenager and senior categories. We concluded that imperfections in face alignment still provide better results than methods using unaligned faces. In future research we will focus more on out-of-plane rotations and high-level classifier fusion to get benefit from individual classifiers.

ACKNOWLEDGMENT

This study is supported by TWIRL (ITEA2 10029 - Twinning Virtual World Online Information with Real-World Data Sources) project.

REFERENCES

[1] J. Bekios-Calfa, J. Buenaposada, and L. Baumela, "Revisiting linear discriminant techniques in gender recognition," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 33, no. 4, pp. 858–864, 2011.

[2] C. Ng, Y. Tay, and B.-M. Goi, "Recognizing human gender in computer vision: A survey," in *PRICAI 2012: Trends in Artificial Intelligence*, ser. Lecture Notes in Computer Science, P. Anthony, M. Ishizuka, and D. Lukose, Eds. Springer Berlin Heidelberg, 2012, vol. 7458, pp. 335–346. [Online]. Available: http://dx.doi.org/10.1007/978-3-642-32695-0_31

[3] P. Phillips, "The FERET database and evaluation procedure for face-recognition algorithms," *Image and Vision Computing*, vol. 16, no. 5, pp. 295–306, Apr. 1998. [Online]. Available: http://dx.doi.org/10.1016/s0262-8856(97)00070-x

[4] G. B. Huang, M. Ramesh, T. Berg, and E. Learned-Miller, "Labeled faces in the wild: A database for studying face recognition in unconstrained environments," University of Massachusetts, Amherst, Tech. Rep. 07-49, October 2007.

[5] A. Gallagher and T. Chen, "Understanding images of groups of people," *2013 IEEE Conference on Computer Vision and Pattern Recognition*, vol. 0, pp. 256–263, 2009.

[6] B. A. Golomb, D. T. Lawrence, and T. J. Sejnowski, "Sexnet: A neural network identifies sex from human faces," in *Proceedings of the 1990 Conference on Advances in Neural Information Processing Systems 3*, ser. NIPS-3. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 1990, pp. 572–577. [Online]. Available: http://dl.acm.org/citation.cfm?id=118850.118953

[7] L. Walawalkar, M. Yeasin, A. Narasimhamurthy, and R. Sharma, "Support vector learning for gender classification using audio and visual cues: A comparison," in *Pattern Recognition with Support Vector Machines*, ser. Lecture Notes in Computer Science, S.-W. Lee and A. Verri, Eds. Springer Berlin Heidelberg, 2002, vol. 2388, pp. 144–159. [Online]. Available: http://dx.doi.org/10.1007/3-540-45665-1_12

[8] C. Shan, "Learning local binary patterns for gender classification on real-world face images," *Pattern Recognition Letters*, vol. 33, no. 4, pp. 431 – 437, 2012, intelligent Multimedia Interactivity. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0167865511001607

[9] M. C. Santana, J. Lorenzo-Navarro, and E. Ramón-Balmaseda, "Improving gender classification accuracy in the wild," in *CIARP (2)*, ser. Lecture Notes in Computer Science, J. Ruiz-Shulcloper and G. S. di Baja, Eds., vol. 8259. Springer, 2013, pp. 270–277.

[10] E. Ramn-Balmaseda, J. Lorenzo-Navarro, and M. Castrilln-Santana, "Gender classification in large databases," in *Progress in Pattern Recognition, Image Analysis, Computer Vision, and Applications*, ser. Lecture Notes in Computer Science, L. Alvarez, M. Mejail, L. Gomez, and J. Jacobo, Eds. Springer Berlin Heidelberg, 2012, vol. 7441, pp. 74–81. [Online]. Available: http://dx.doi.org/10.1007/978-3-642-33275-3_9

[11] P. Dago-Casas, D. Gonzalez-Jimenez, L. L. Yu, and J. Alba-Castro, "Single and cross-database benchmarks for gender classification under unconstrained settings," in *Computer Vision Workshops (ICCV Workshops), 2011 IEEE International Conference on*, 2011, pp. 2152–2159.

[12] T. Ojala, M. Pietikainen, and T. Maenpaa, "Multiresolution gray-scale and rotation invariant texture classification with local binary patterns," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 24, no. 7, pp. 971–987, 2002.

[13] E. Makinen and R. Raisamo, "Evaluation of gender classification methods with automatically detected and aligned faces," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 30, no. 3, pp. 541–547, 2008.

[14] http://mplab.ucsd.edu, "The MPLab GENKI Database, GENKI-4K Subset," 2011.

[15] P. Viola and M. J. Jones, "Robust real-time face detection," *International Journal of Computer Vision*, vol. 57, pp. 137–154, 2004, 10.1023/B:VISI.0000013087.49260.fb. [Online]. Available: http://dx.doi.org/10.1023/B:VISI.0000013087.49260.fb

[16] G. Bradski, *Dr. Dobb's Journal of Software Tools*.

[17] H. Rowley, S. Baluja, and T. Kanade, "Neural network-based face detection," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 20, no. 1, pp. 23 –38, jan 1998.

[18] S. Milborrow and F. Nicolls, "Locating facial features with an extended active shape model," in *Computer Vision ECCV 2008*, ser. Lecture Notes in Computer Science, D. Forsyth, P. Torr, and A. Zisserman,

Eds. Springer Berlin / Heidelberg, 2008, vol. 5305, pp. 504–513. [Online]. Available: http://dx.doi.org/10.1007/978-3-540-88693-8_37

[19] A. B. A. Graf and S. Borer, "Normalization in support vector machines," in *DAGM-Symposium*, ser. Lecture Notes in Computer Science, B. Radig and S. Florczyk, Eds., vol. 2191. Springer, 2001, pp. 277–282.

[20] C.-C. Chang and C.-J. Lin, "Libsvm: A library for support vector machines," *ACM Trans. Intell. Syst. Technol.*, vol. 2, no. 3, pp. 27:1–27:27, May 2011. [Online]. Available: http://doi.acm.org/10.1145/1961189.1961199

[21] E. Makinen and R. Raisamo, "An experimental comparison of gender classification methods," *Pattern Recognition Letters*, vol. 29, no. 10, pp. 1544 – 1556, 2008. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0167865508001116

[22] G. Guo, C. Dyer, Y. Fu, and T. Huang, "Is gender recognition affected by age?" in *Computer Vision Workshops (ICCV Workshops), 2009 IEEE 12th International Conference on*, 2009, pp. 2032–2039.

[23] J. Bekios-Calfa, J. M. Buenaposada, and L. Baumela, "Robust gender recognition by exploiting facial attributes dependencies," *Pattern Recognition Letters*, vol. 36, no. 0, pp. 228 – 234, 2014. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0167865513001864

[24] G. Tsoumakas, I. Katakis, and I. P. Vlahavas, "Random k-labelsets for multilabel classification," *IEEE Trans. Knowl. Data Eng.*, vol. 23, no. 7, pp. 1079–1089, 2011.

[25] K. Ricanek and 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*, 2006, pp. 341–345.