Children Gender Recognition Under Unconstrained Conditions Based on Contextual Information

Riccardo Satta, Javier Galbally and Laurent Beslay
Joint Research Centre, European Commission, Ispra, Italy
Email: {riccardo.satta,javier.galbally,laurent.beslay}@jrc.ec.europa.eu

Abstract*—One of the biggest challenges faced by law enforcement entities in the present digital era, is fighting against on-line Child Sexual Abuse (CSA), due in particular to the massive amount of data that they receive for analysis. Pattern recognition system can provide an aid, e.g., to ease the identification of both the perpetrator and the victim of the crime. In particular, ancillary cues related the identity of the involved persons, like age, race or gender, can represent a significant aid for identification. These cues can be estimated using statistical classifiers on face features. In this work, we explore one of these ancillary cues, namely the gender. The research community has provided methods for gender recognition able to achieve good performance with adults. However, in the case of CSA, victims are minors (typically, very young children). Children gender recognition may be difficult even for humans, due to the lack of many gender-specific face traits usually present in adult faces. Totally uncontrolled poses and illumination conditions, that might be found in CSA material, represent an additional issue. We propose to tackle this problem by the use of *contextual* information to complement face features used by traditional algorithms. In particular, we exploit the image context of the face, that is, the portion of the image surrounding the face. This is motivated by the usage that humans themselves make of face external information, such as the hair or earrings, to take decisions on this task. The proposed approach is tested on a novel data base of faces of children, collected from royalty-free stock-photography web sites, which show totally unconstrained conditions. The reported results are promising and set the way for a deeper study of the use of the face context for estimating ancillary identification cues.

I. Introduction

The digital era has lead to a real breakthrough in the field of multimedia data. With the development of Internet and the new generation of smart portable devices, the acquisition, storage, processing and sharing of images and video is within everybody's reach. However, this reality, which constitutes an undeniably positive development in our everyday life, can also become a very difficult issue to address under specific scenarios where the acquired multimedia data represents in itself a crime. This is the case for example of images and videos related to child sexual abuse (CSA).

The fact that nowadays the digital technology is accessible to all has made fairly easy for child offenders to capture and share illegal content of their criminal actions. As such, one of the biggest challenges faced by law enforcement entities from all over the world in the fight against on-line child sexual abuse, is to manage in an efficient manner the massive amount of data that they receive for analysis almost on a daily basis.

Before the advent of the digital revolution, forensic analysis of this type of data was carried out manually. However, in the current context, with thousands of stored terabytes, such an approach is no longer feasible. Consulted forensic experts from different European law enforcement bodies (e.g., Europol, France, Italy or Spain) have highlighted the need for specific tools that can help them to i) automatically classify between legitimate and illegal content, and ii) process the potentially criminal data, to identify both perpetrator and victim of the crime. Such aiding applications would have a deep impact in the analysis of the two-year delayed stock data related to CSA cases.

Face detection and recognition plays a major role in this very challenging classification and identification task. Unfortunately, in most cases, due to the lack of enrolled data (i.e., images linked to a known identity), direct identification of subjects that appear in images and videos is not possible. However, many other ancillary cues directly related to face can also represent a notable aid for narrowing down the search in a particular crime. This is the case, for instance, of a person's age, race, or gender. In particular, it is especially relevant to focus the feature extraction process on victims, as they are generally more exposed and visible than criminals (who tend to stay out of the scene). In CSA cases, victims are minors ranging from teenagers to very young children, therefore, for this distinct scenario it gains a lot of importance the development of automatic pattern recognition algorithms adapted to the specificities of kids.

In the present article we explore one of the face-ancillary information related to CSA data analysis, which to the authors' best knowledge, has still not been exploited in most of the few aiding-forensic tools available in the market. In particular, we propose a new algorithm for children gender estimation under totally uncontrolled conditions such as the ones that might be expected to be found in the specified CSA forensic scenario. The problem addressed presents two major challenges: i) on the one hand, as already mentioned above, the type of data involved. CSA related content is collected under a totally unconstrained environment. This includes: different types of acquisition devices, illumination, background, occlusions, movement, angle changes, pose, etc. ii) On the other hand, the particular difficulties linked to children gender estimation. In general, children do not present most of the gender distinctive features that are developed with age, so that in many occasions it is difficult even for humans to distinguish the gender of a kid in the absence of any information other than the face. This issue becomes more acute at early age.

Due to the two major problems described above, the accuracy of state-of-the-art gender recognition algorithms drops significantly when it is tested on CSA-like data. To overcome such a performance decrease, we propose the use of *contextual* information to complement face features used by traditional algorithms (usually tested on adult images acquired under controlled conditions). In particular, we exploit the image context of the face, that is, the portion of the image surrounding the face. The use of the image context is mainly motivated by the utilisation that humans themselves make of face external information, such as the hair or earrings, to take decisions on this task [1].

Following the very successful example of the Labeled Faces in the Wild DB [2], the novel approach has been tested on a database of royalty-free kid images taken in real life situations and downloaded from publicly accessible web picture repositories. This way, the dataset emulates, as far as possible, the totally uncontrolled conditions of CSA related data. The proposed method, exploiting the fusion of face-features with different context-based characteristics, has shown some encouraging improvements with respect to the performance of state-of-the-art algorithms working only with facial information.

The rest of the article is structured as follows. First, in Sect. II we analyse the current state-of-the-art on techniques for gender recognition, and point out some relevant studies on the specific task of children gender recognition. We then describe the proposed gender recognition approach based on the combination of face and contextual information, in Sect. III, and test it in a novel database of children images as described in Sect. IV. The experimental results are reported in Sect. V. Finally, Sect. VI sums up the work and concludes the paper.

II. RELATED WORKS

Various gender estimation techniques have been proposed in the literature to infer the gender of a person from his face image. Usually, these methods assume that a face has already been detected and segmented. The processing pipeline then resembles the one of a typical machine learning problem: i) Image representation: a set of features is extracted from the image; ii) Classification: a statistical classifier is used to estimate the gender.

Image representation plays an important role in gender estimation. The simplest kind of representation used is a matrix of the raw pixel values [3], [4], possibly associated with feature selection methods to reduce the space dimensionality (e.g., PCA [5], LDA [6]). Although it generally leads to a good classification performance, such representation is not robust to misalignments and to pose changes. More clever representations that can increase the level of robustness include Haar-like features [7] and Local Binary Pattern (LBP) [8], [9].

The classification function is normally learnt from a training set of labelled images (male face/female face). The most widely used classifiers to this aim are AdaBoost [10], [7] and Support Vector Machines (SVM) [9], [4], [3].

Gender estimation techniques address the same challenges as face recognition, namely pose and facial expression variations, non-uniformity of the background, partial occlusions and light variations. Unfortunately, much research has been conducted using data sets of images acquired under controlled conditions, with only or mostly frontal poses, no occlusions, consistent lighting and a clean background (e.g., the widely used FERET benchmark data set [11]). As a result, such methods fail to achieve a reasonable performance when used in real-life images. This situation is changing since the introduction of the Labeled Faces in the Wild (LFW) dataset[1] [2], a set of labelled face images taken in unconstrained environments from real-life pictures, which is also used for benchmarking face recognition techniques. In fact, some recent gender recognition techniques used LFW in order to better design systems and benchmark them [12], [8]. Current state-of-the art techniques are able to achieve a classification performance of about 94% correctly labelled samples in such data set [8]. The drawback of the LFW-DB for the scenario considered in the present work is that it mainly contains images of adults, and therefore is not suitable for children gender recognition assessment.

It is particularly relevant, for the purposes of this work, to understand if and how age can affect the performance of gender estimation techniques. In fact, the face of young children usually does not show many of the typical gender traits, so that gender estimation of kids can be a difficult task even for humans, especially in absence of gender-stereotyped cues (e.g. the hairstyle) [1]. Unfortunately, this aspect has been largely overlooked in the scientific literature. Benabdelkader and Griffin [13] reported that a disproportionately large number of elderly females and young males were misclassified by their method when tested on a data set of 12,964 face images. Similar results are shown in [14], where the performance of a gender estimation method that exploits facial attributes dependencies has been evaluated with respect to the age range.

A more comprehensive study has been conducted by Guo *et al.* [15] on a set of 8,000 faces ranging from 0 to 93 years old. They showed that the classification accuracy in the age range 0-19 can be up to 13% lower compared to the age range 20-60. The best performance obtained was about 97% for the adult age range (20-60) but below 88% for the young age range (0-19). It is worth pointing out that in that work, the age range 0-19 mostly included minors belonging to the second half of the range (10-19), for whom gender estimation is likely to be easier, but significantly fewer young children (below 10), that possibly are the most difficult cases.

To the best of our knowledge, up to now no one has assessed the performance of gender estimation techniques on a database of faces of young children, nor proposed specific algorithms for this kind of images. In fact, as shown in the previous studies, the visual appearance of the face may not be enough to obtain a high gender recognition accuracy in such cases.

In the present work this problem is addressed for totally unconstrained acquisition conditions by combining visual face-specific features with “ancillary” (i.e., contextual) information related to the presence of gender-stereotyped cues (i.e., hairstyle or clothes), just as humans do [1]. It is worth pointing out that, although contextual features and attributes have been already exploited in combination with facial features in different tasks (e.g., face verification [16]), to our best knowledge they have not been used for gender recognition children.

[1] http://vis-www.cs.umass.edu/lfw/index.html

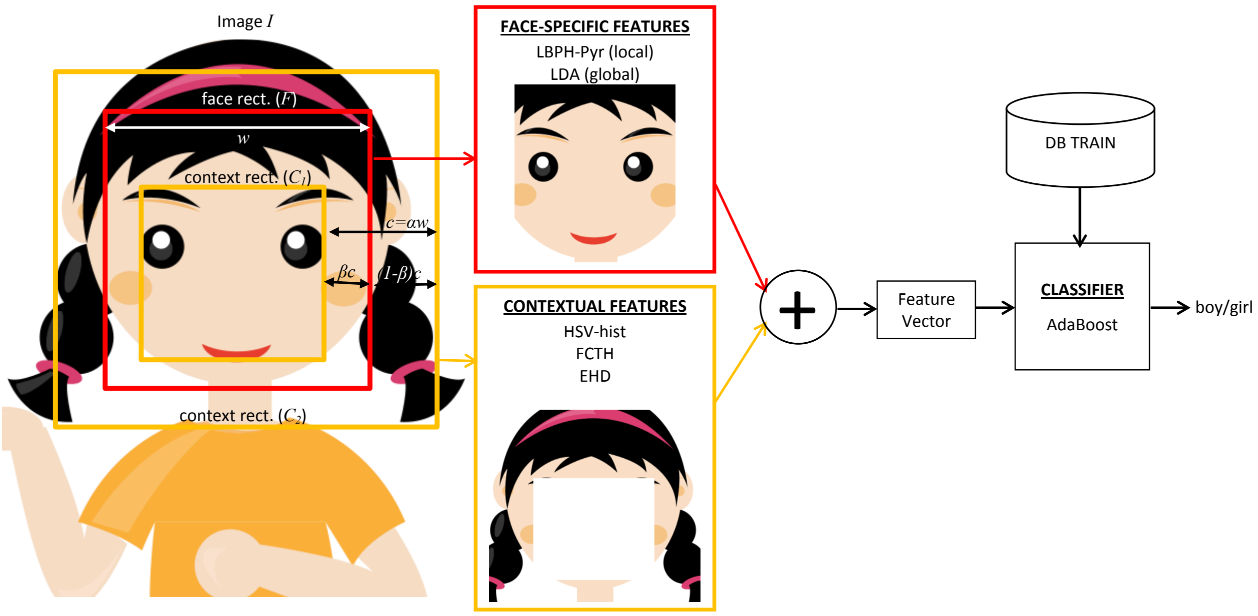

Fig. 1. General diagram of the proposed method for children gender recognition based on the fusion of face-specific and contextual features. A cartoon appears instead of the image of a real kid due to data protection constraints, that do not permit the public exposure of any image that could imply a link (not even fake as in this case) between kids and CSA content.

III. CHILDREN GENDER RECOGNITION BASED ON CONTEXT FEATURES

In this Section, we describe the proposed approach to extract contextual information from the face image and combine it with traditional face features for gender recognition.

Let I be the image containing a face, and F the bounding box containing the face (see Fig. 1). We define the *face context* as the region of the image contained between two rectangles C_1 and C_2, $C_1 < C_2$, both centered in F. The *width* of the context region (denoted as c in Fig. 1), that is, the space between the two rectangles, is proportional to the width w of F: $c = \alpha \cdot w$.

The distance between C_1 and F perimeters, and between F and C_2 perimeters, are proportional to c and respectively equal to $\beta \cdot c$ and $(1 - \beta) \cdot c$. The two parameters α and β therefore control respectively the *size* of the face context, and how much it overlaps with the face rectangle.

Once the face and its context have been segmented, we first extract face-specific features from the face, as in other gender recognition techniques (see Sect. II). Separately, we characterise the context, using an *image descriptor*, as those used for scene recognition and image retrieval tasks. The two sets of data (face-specific and context-based) are then combined at feature level, by concatenating the respective feature vectors into a single one. A statistical classifier, previously trained on a training data set of labelled samples, is finally used to recognise the gender.

As face-specific features (extracted only from F), one *local* feature set and one *global* feature set are used in the work:

- *Local face-specific feature set*. It is a block-based, multi-scale modification of the approach presented in [8] for adult gender recognition under unconstrained condition based on the very popular Local Binary Patterns. The face image F is first normalized to a square of $64{\times}64$ pixels and converted to grey F_n. Then, a sliding window of $8{\times}8$ pixels is moved through the image, with a step of 4 pixels so that consecutive windows are partly overlapped. An LBP Histogram is computed in each window position; histograms are then concatenated into a single feature vector. The same approach is repeated at two other scales, by resizing the image to a square of $32{\times}32$ pixels, then to $16{\times}16$ pixels. We refer to this feature set as LBPH-Pyr in the following. By using different sub-blocks of the image, and different scales, LBPH-Pyr should better capture localised gender-specific face characteristics.
- *Global face-specific feature set*. It is the classical LDA representation widely used in face recognition also referred to in some cases as *fisherfaces* [17], and exploited for gender recognition for instance in [6]. The original $64 \times 64 = 4,096$ dimensions contained in the grey-scaled normalized face image (F_n) are reduced to a 400 dimensional LDA-space (computed from a set of training data as described in Sect. IV).

As for the contextual features, we chose three scene representation descriptors widely used in the literature:

- The colour histogram in the HSV colour space (HSV-Hist), with respectively 32, 24 and 4 bins (the V channel being the most under-sampled as it encodes mainly information on illumination).
- The Fuzzy Colour and Texture Histogram (FCTH) descriptor [18], that encodes both colour and texture information.
- The Edge Histogram Descriptor (EHD), defined in the

MPEG-7 standard [19], which encodes information on textures and edges.

Finally, we used the OpenCV [20] implementation of Real AdaBoost, as classifier, with 2000 weak learners.

IV. Database and Experimental Protocol

In the field of child sexual abuse, in addition to the technical difficulties already described (i.e., unconstrained acquisition conditions and age of the victims), researchers have to cope with the challenge of the criminal nature of data. As such, it constitutes very sensitive information only available for forensic examiners within law enforcement facilities. Data protection laws do not allow the disclosure or distribution of such data independently of the final use it will be given.

Therefore, the development of automatic algorithms for the analysis of this type of data has to be carried out on datasets imitating, to some extent, the most relevant characteristics found in real CSA content. Unfortunately, as mentioned in Sect. II, currently there is no publicly available database that fits the goals of the present work: young children gender recognition under totally unconstrained conditions.

Consequently, following the very successful example of Labelled Faces in the Wild [2] and of some other valuable works in adult gender recognition [21], for the current study we have constructed a new database using royalty-free images obtained from publicly accessible on-line photo repositories, the "Gender-Labelled Children Faces in the Wild Database" (GLCFW-DB). The dataset generation process followed three successive steps:

- Step 1: Download. In total, 4,270 real-life kid-related images were downloaded from: flickr (736), freedigitalphotos (1,520) and freepik (2,014). The images have been retrieved by searching the queries "kid", "child", "children", "boy" and "girl" on these web sites.
- Step 2: Face detection. The Viola-Jones detector was used for this purpose. Images could contain none, one or several faces. The detector was set to generate many false positives in order to miss as few faces as possible.
- Step 3: Manual gender tagging. Segmented images were labelled as masculine, feminine or discarded to generate the database ground-truth. A detected face was discarded if: it did not contain a face (detection error), it contained the face of an adult, the age was unclear to determine if the subject was a kid or already a teenager, or the gender could not be established (due to the difficulty of distinguishing between young baby-boys and baby-girls based only of facial features).
- Step 4: Face normalization. All faces were normalized in size to grey-scale images of 64×64 pixels (F_n) to extract the face-specific set of features (i.e., LBPH-Pyr and LDA as described in Sect. III).

The final GLCFW-DB contains a total 1,840 face images of children, corresponding to 876 boys (48%) and 964 girls (52%). The database poses a real challenge for automatic gender recognition as it comprises real-like images of young children collected with no constraints in terms of: illumination, pose, occlusions (even make-up), race, age, resolution of the images, acquisition device, etc. Unfortunately, due to data protection constraints, no real images from the database can be publicly exposed in order to avoid creating any implicit link between the kids that appear in the pictures and CSA content. Therefore, no database samples can be shown in the article as this could drive uninformed readers to create wrong assumptions.

FACE ONLY

	1st rep.	2nd rep.	3rd rep.	4th rep.	5th rep.	average
LBPH-Pyr	69.42%	66.49%	68.44%	71.06%	70.08%	**69.10%**
LDA	55.39%	52.88%	56.91%	56.26%	53.97%	55.08%

CONTEXT ONLY

	1st rep.	2nd rep.	3rd rep.	4th rep.	5th rep.	average
HSV-Hist	59.87%	63.98%	61.91%	63.53%	63.00%	**62.46%**
FCTH	58.11%	58.10%	57.89%	54.62%	58.87%	57.52%
EHD	55.26%	56.15%	52.80%	56.47%	57.02%	55.54%

TABLE I. Performance on the GLCFW database, in terms of Classification Accuracy over a 5-fold validation, using face-only and context-only gender recognition methods. The best average performance is highlighted in bold.

FACE + CONTEXT

	1st rep.	2nd rep.	3rd rep.	4th rep.	5th rep.	average
LBPH-Pyr + HSV-Hist	76.61%	74.76%	74.86%	76.61%	72.91%	**74.89%**
LBPH-Pyr + FCTH	73.67%	72.69%	75.08%	74.97%	73.01%	73.88%
LBPH-Pyr + EHD	73.45%	73.45%	72.90%	74.86%	72.58%	73.45%
LDA + HSV-Hist	57.02%	59.74%	58.54%	59.63%	57.13%	58.41%
LDA + FCTH	56.69%	54.73%	58.32%	55.71%	55.28%	56.14%
LDA + EHD	54.84%	53.75%	56.15%	54.95%	55.82%	55.11%

TABLE II. Performance on the GLCFW data base, in terms of Classification Accuracy over a 5-fold validation, combining face-specific with context features. The best average performance is highlighted in bold.

As a result of the process followed to generate the database, one of its limitations is the lack of any age information from the children. Therefore, it does not allow to perform experiments related to the accuracy of gender recognition algorithms with respect to age.

In the experiments, a 5-fold validation strategy was followed. The training and test sets were randomly selected for each fold and contained half of the boys and girls images (i.e., 438 boys and 482 girls). The training set is used to compute the LDA transformation matrix and to tune the AdaBoost classifier. Results are reported in terms of average Classification Accuracy, that is, the percentage of correctly classified samples in the test set.

The database is available under email request to the authors only for research purposes. The distributed data consists of: the original downloaded images, the coordinates of the detected faces for each image, the ground-truth tags for each detected face (i.e., masculine, feminine or discarded), and the faces used for training and test in the 5-fold validation protocol. This way, following the path opened by the LFW-DB, results can be fairly and objectively compared with other approaches proposed in the future, so that the evolution of the state-of-the-art in this challenging task (i.e., children gender recognition under unconstrained conditions) may be established.

V. RESULTS

Table I-top shows the performance attained using only the face-specific features LBPH-Pyr and LDA. The first five columns present the classification accuracy obtained in each repetition of the 5-fold validation, while the last column give the average classification accuracy over the 5 repetitions. The parameters α and β have been empirically set to 0.4 and 0.3,respectively. As expected, such performance is low, due to the combined effect of $i)$ little difference of face traits of young males and females, and $ii)$ completely unconstrained poses and illumination conditions of the images in the GLCFW database. In particular, the classification accuracy of LDA is in practice equivalent to random guess. LBPH-Pyr performs better, still an accuracy of 69.10% is quite unsatisfactory for a practical usage. The difference between LDA and LBPH-Pyr is mainly motivated by the fact that LDA looks globally at the raw pixels, and therefore is much more sensitive to pose and illumination variations. LBPH-Pyr instead, being capable to distinguish localised face traits, shows more robustness to such variations.

The context, when used alone (see Table I-bottom), also leads to poor classification accuracies. Colour information seems more relevant to characterise the context; in fact, the worst performing context descriptor proves to be EHD, which encodes only information on textures and edges. This is most probably due to the fact that color is an efficient way of encoding information related to long hair, which represents one of the most typical contextual cues used for gender estimation.

In Table II, we show the performance attained when face and context are combined. LBPH-Pyr receives a notable boost of almost 6% (from 69.10% to 74.89%) when combined with HSV-Hist. The combination of LBPH-Pyr with the other context descriptors also improves the performance with respect to LBPH-Pyr alone. LDA also performs better when combined with context descriptors, although the final performance remains below 60%. Once again, this confirms that LDA is not suitable for gender classification under unconstrained conditions.

VI. CONCLUSION

In this work, we proposed a novel approach for children gender recognition under unconstrained conditions. The method exploits not only face-specific information, which in the case of children is in many cases insufficient, but also contextual features designed to model gender-stereotyped cues such as the presence of long hair or earrings. The method has been tested on a new database of real-like children images which has been made public for future comparison with other algorithms. Results show that, the addition of face-ancillary information can help to improve the accuracy of traditional face-based gender recognition algorithms, under this very challenging conditions.

The proposed technique can provide a significant aid in the field of fight against child sexual abuse, where the identification of the victim can lead not only to the offender but may also help in other cases related, for instance, to missing children.

Although the achieved experimental results are quite promising they are far from the performance rates obtained on adult data acquired under controlled conditions. Still much research is needed in the area of children gender recognition under realistic conditions for its practical use as a supporting tool for forensic examiners, who are currently overwhelmed by the amount of data related to CSA cases. Unfortunately, the range of such type of aiding-forensic tools in the market is still quite reduced and, in many cases, they do not exploit all the information that could potentially be automatically extracted from the seized material. In this context, pattern recognition algorithms such as the one described in the present work can contribute in a significant manner to further develop the existing applications, by adding complementary search capabilities.

REFERENCES

[1] H. A. Wild, S. E. Barrett, M. J. Spence, A. J. OToole, Y. D. Cheng, and J. Brooke, "Recognition and sex categorization of adults and childrens faces: Examining performance in the absence of sex-stereotyped cues," *Journal of Experimental Child Psychology*, vol. 77, pp. 269–291, 2000.

[2] 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, Tech. Rep., 2007.

[3] B. Moghaddam and M.-H. Yang, "Learning gender with support faces," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 24, pp. 707–711, 2002.

[4] E. Makinen and R. Raisamo, "Evaluation of gender classification methods with automatically detected and aligned faces," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 30, pp. 541–547, 2008.

[5] K. Balci and V. Atalay, "PCA for gender estimation: Which eigenvectors contribute?" in *Proc. IEEE Int. Conf. on Pattern Recognition (ICPR)*, 2002, p. 30363.

[6] J. Bekios-Calfa, J. M. Buenaposada, and L. Baumela, "Revisiting linear discriminant techniques in gender recognition," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 33, pp. 858–864, 2011.

[7] G. Shakhnarovich, P. A. Viola, and B. Moghaddam, "A unified learning framework for real time face detection and classification," in *Proc. IEEE Int. Conf. on Automatic Face and Gesture Recognition (ICFGR)*, 2002.

[8] C. Shan, "Learning local binary patterns for gender classification on real-world face images," *Pattern Recognition Letters*, vol. 33, pp. 431–437, 2012.

[9] A. Hadid and M. Pietikainen, "Combining appearance and motion for face and gender recognition from videos," *Pattern Recognition*, vol. 42, pp. 2818–2827, 2009.

[10] S. Baluja and H. A. Rowley, "Boosting sex identification performance," *Int. Journal of Computer Vision*, vol. 71, pp. 111–119, 2007.

[11] P. J. Phillips, H. Moon, S. A. Rizvi, and P. J. Rauss, "The FERET evaluation methodology for face-recognition algorithms," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 22, pp. 1090–1104, 2000.

[12] 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 *Proc. IEEE Int. Conf. on Computer Vision Workshops (ICCV-W)*, 2011, p. 21522159.

[13] C. BenAbdelkader and P. Griffin, "A local region-based approach to gender classification from face images," in *Proc. IEEE Conf. on Computer Vision and Pattern Recognition Workshops (CVPR-W)*, 2005.

[14] J. Bekios-Calfa, J. M. Buenaposada, and L. Baumela, "Robust gender recognition by exploiting facial attributes dependencies," *Pattern Recognition Letters*, 2013.

[15] G. Guo, C. R. Dyer, Y. Fu, and T. S. Huang, "Is gender recognition affected by age?" in *Proc. IEEE Int. Conf. on Computer Vision Workshops (ICCV-W)*, 2009, p. 20322039.

[16] N. Kumar, A. C. Berg, P. N. Belhumeur, and S. K. Nayar, "Attribute and simile classifiers for face verification," in *Proc. IEEE International Conf. on Computer Vision (ICCV)*, 2009.

[17] P. Belhumeur, J. Hespanha, and D. Kriegman, "Eigenfaces vs fisherfaces: Recognition using class specific linear projection," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 19, pp. 711–720, 1997.

[18] S. A. Chatzichristofis and Y. S. Boutalis, "Fcth: Fuzzy color and texture histogram - a low level feature for accurate image retrieval," in *Proc. 9th Int. Workshop on Image Analysis for Multimedia Interactive Services (WIAMIS)*, may 2008, pp. 191–196.

[19] T. Sikora, "The mpeg-7 visual standard for content description - an overview," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 11, no. 6, pp. 696–702, Jun. 2001.

[20] G. R. Bradski, "The OpenCV Library," *Dr. Dobb's Journal of Software Tools*, 2000.

[21] S. Y. D. Hu, B. Jou, A. Jaech, and M. Savvides, "Fusion of region-based representations for gender identification," in *Proc. IEEE Int. Joint Conf. on Biometrics (IJCB)*, 2011.