University of Groningen

Gender recognition from face images using trainable shape and colour features

Azzopardi, George; Foggia, Pasquale; Greco, Antonio; Saggese, Alessia; Vento, Mario

Published in:
International Conference of Pattern Recognition

DOI:
[10.1109/ICPR.2018.8545771](https://doi.org/10.1109/ICPR.2018.8545771)

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version
Publisher's PDF, also known as Version of record

Publication date:
2018

[Link to publication in University of Groningen/UMCG research database](#)

Citation for published version (APA):

Azzopardi, G., Foggia, P., Greco, A., Saggese, A., & Vento, M. (2018). Gender recognition from face images using trainable shape and colour features. In *International Conference of Pattern Recognition* (pp. 1983-1988). IEEE. <https://doi.org/10.1109/ICPR.2018.8545771>

Copyright

Other than for strictly personal use, it is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license (like Creative Commons).

The publication may also be distributed here under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license. More information can be found on the University of Groningen website: <https://www.rug.nl/library/open-access/self-archiving-pure/taverne-amendment>.

Take-down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Downloaded from the University of Groningen/UMCG research database (Pure): <http://www.rug.nl/research/portal>. For technical reasons the number of authors shown on this cover page is limited to 10 maximum.

Gender recognition from face images using trainable shape and color features

George Azzopardi

Faculty of Science and Engineering

Johann Bernoulli Inst. for Math. and CompSc.

University of Groningen, the Netherlands

Email: g.azzopardi@rug.nl

Pasquale Foggia, Antonio Greco, Alessia Saggese, Mario Vento

Faculty of Engineering

Dept. of Computer Eng. and Electrical Eng. and Applied Mathematics

University of Salerno, Italy

Email: agreco, asaggese, pfoggia, mvento@unisa.it

Abstract—Gender recognition from face images is an important application and it is still an open computer vision problem, even though it is something trivial from the human visual system. Variations in pose, lighting, and expression are few of the problems that make such an application challenging for a computer system. Neurophysiological studies demonstrate that the human brain is able to distinguish men and women also in absence of external cues, by analyzing the shape of specific parts of the face. In this paper, we describe an automatic procedure that combines trainable shape and color features for gender classification. In particular the proposed method fuses edge-based and color-blob-based features by means of trainable COSFIRE filters. The former types of feature are able to extract information about the shape of a face whereas the latter extract information about shades of colors in different parts of the face. We use these two sets of features to create a stacked classification SVM model and demonstrate its effectiveness on the GENDER-COLOR-FERET dataset, where we achieve an accuracy of 96.4%.

I. INTRODUCTION

The exponential growth of the advanced video analysis technologies is raising more and more interest for systems that process real sequences of images and provide large amounts of intelligent data about the considered scenario. The commercial interest for such solutions mainly occurs in the retail field, where the characterization of the customers visiting a store gives remarkable insights to carry out extensive market analysis. Gender recognition from face images or video frames is one of the most required smart applications, since the knowledge of this biometric feature allows the retailers to customize the promotional material and the placement of the products in their shops. Although gender identification seems to be a trivial task for a human, the scientific community is still active in this field of research. Indeed, in recent years several researchers attempted to understand how humans identify the gender of a person so as to propose novel gender recognition methods [1].

The difficulties in the identification of the gender from face images are mainly related to two factors. First, the capture process may cause challenging variations on the face picture due to different poses, illumination conditions and occlusions. Second, there are intrinsic differences between human faces, both in terms of genetic traits (facial features, age, race) and occasional variations (expression) or accessories (glasses, sunglasses, scarves, hats, earrings).

The former problem is generally addressed by performing various pre-processing steps, which have the aim to normalize the face image with respect to the possible variations. The most typical normalization procedure is the rotation of the face image in order to horizontally align the pose by using the position of the eyes as reference. The latter seems to be a more challenging problem. There are some obvious indicators that allow to easily distinguish in most of the cases men and women. For example, the presence of beard or facial hairs are clear clues for recognizing a man, while make-up is a more discriminative feature of women. Other discriminative elements for a woman, not as effective as the previous ones, may be long hair and earrings. Such preliminary observations suggest two questions: (i) how to detect such strong indicators? (ii) which features does a human use to distinguish men and women in absence of obvious cues?

Considering these observations, one may suggest that a face descriptor for gender recognition may include texture features, which seem to be suitable for detecting features such as beard and rough skin. Indeed, several methods that use the histograms of LBP features [2], [3], [4], [5] have been proposed in the last decade. Texture features are, however, not able to deal with more challenging situations. For example, such descriptors may not have elements to distinguish young men without beard and women with short hair, or to discern elderly women with rougher skin due to the wrinkles and adult men.

Other researchers use directly the pixel intensity values [6], [7], [8], assuming that the classifiers have the capability to autonomously identify the most discriminant parts of the face for gender recognition. However such methods are not invariant to translation, even using face alignment algorithm to normalize the pose, since a shift of a few pixels may completely change the resulting feature vector.

Physiological studies demonstrate that men and women have different traits due to the effect of testosterone and estrogen on the face [9]. Indeed, a man has a wide jaw, pronounced chin and cheekbones and smaller eyes, while the estrogen bring a greater amount of subcutaneous fat and make the traits of a woman softer. The faces appear to the brain more or less masculine or feminine according to the ratio between

Fig. 1: Architecture of the proposed method.

the two types of hormones, but these differences are not so easy to recognize for an automatic algorithm. Taking inspiration from these studies, some scientists use histograms of gradients [10] to model descriptors able to encode the shapes of human faces [11], [12] or fuse texture, color and shape features to capture all the different aspects in a single classifier [13], [14], [15]. Similarly, other studies propose the extraction of local features from particular parts of the face [16], [17], namely the facial landmarks [18], so as to capture such soft differences.

In the last years we have seen the benefit of using trainable feature detectors in contrast to handcrafted ones. Convolutional neural networks have demonstrated their effectiveness in various applications [19], [20], [21], [22]. Similarly, COSFIRE features, which can be easily interpreted, have also shown their strength in a number of applications [23]. In particular, we already demonstrated the effectiveness of edge-based COSFIRE features for gender recognition in face images [24] [25]. A new type of COSFIRE features have been introduced recently [26], which encode the mutual spatial arrangement of color blobs. These are especially beneficial in images with diffuse boundaries such as faces.

In this paper we propose an algorithm inspired by the cognitive process which enables a human to recognize the gender of a person by processing and combining the observations of specific facial features. The idea is to fuse trainable shape and color features by using the COSFIRE filters described in [23] and [26], in order to encode (1) the physiological differences between male and female somatic traits and (2) the obvious cues related to the presence of makeup, especially in the areas of eyes and lips. The main advantage of this choice is that the edge-based and color-blob-based COSFIRE features are learned from the training data. To the best of our knowledge, there are no papers in the literature which address the problem of gender recognition by defining an automatic procedure for selecting and combining discriminative facial features related to intrinsic human differences and external cues.

II. METHOD

A. Overview

We combine two types of trainable COSFIRE filters for gender recognition from face images. One type being the edge-based COSFIRE filters, whose effectiveness has already been demonstrated in [24]. They take as input the responses of some orientation-selective Gabor filters at certain positions and combine them with a nonlinear function. The other type are called color-blob-based COSFIRE filters and are highly effective to detect patterns with diffused edges and of different colors. The fact that faces are full of sharp (e.g. eyes) and diffuse boundaries (e.g. cheek bones) also with different shades of colors, our hypothesis is that these two types of filters can be used to extract features that are complementary to each other.

Figure 1 depicts a high level architecture of the method that we propose.

B. Face detection and normalization

We apply the highly robust and scale-invariant Viola-Jones algorithm [27] to detect the faces in given images. Then we resize the detected faces to a fixed size of 128×128 pixels.

C. Edge-based COSFIRE filters

Edge-based COSFIRE filters are the original type of COSFIRE filters that were published in [23]. They are trainable and they use as input the responses of some Gabor filters. The term trainable refers to the ability to configure a COSFIRE filter whose selectivity is automatically determined by the automatic analysis of a single training example, usually referred to as prototype. The configuration procedure consists of three steps, namely *filter-detection-description*. First, we apply a bank of asymmetric Gabor filters (with 16 orientations and one scale) to a given prototype, threshold their responses with a fraction of 0.1 of the global maximum Gabor response, and superimpose the 16 thresholded response maps. Second, we consider a set of concentric circles with given radii around the center of the given prototype, and choose the locations at which we achieve local maxima along the circles. In the description stage

Fig. 2: Examples of face detection on GENDER-COLOR-FERET images.

we form a set of 4-tuples $S = \{(\lambda_i, \theta_i, \rho_i, \phi_i) \mid i = 1 \dots n\}$ where each tuple i describes four properties of a detected point along the concerned circles; λ_i is the scale and θ_i is the orientation of the Gabor filter that gives the maximum response at the i -th local maximum point, while ρ_i and ϕ_i are the distance and polar angle with respect to the center, respectively.

The application of an edge-based COSFIRE filter requires three operations, namely *filter-blur-shift*, for each tuple and one operation, *multiply*, that combines the intermediate resulting maps from the tuples. Since the operations on the tuples are independent of each other, parallel processing could be used to process the tuples simultaneously. For each tuple i , we first *filter* a given image with a Gabor function of scale λ_i and orientation θ_i . In order to allow for some tolerance with respect to the preferred position we then *blur* the Gabor response map with a blurring function whose standard deviation σ is a linear function of the distance ρ_i from the center of the filter: $\sigma = \sigma_0 + \alpha\rho_i$. In this work we set the constant parameters $\sigma_0 = 0.67$ and $\alpha = 0.1$ as suggested in [24]. Subsequently, we shift the blurred Gabor responses by the vector $(\rho_i, 2\pi - \phi_i)$ so that all involved Gabor responses meet at the support center of the COSFIRE filter. Finally, we combine all the blurred and shifted Gabor responses with geometric mean, which is essentially multiplication. More technical details can be found in [23].

D. Color-blob-based COSFIRE filters

Color-blob-based COSFIRE filters are conceptually similar to the above type of filters. There are three main differences. Firstly, instead of using the orientation-selective Gabor filters we use Difference-of-Gaussians (DoG) filters, which are suitable for blob detection. Secondly, we use the Lab color space and configure tuples from each color channel. Thirdly, we use a different configuration approach to identify points of interest within a given prototype. In the edge-based approach we use a set of concentric circles and find local maximum Gabor responses along such circles. For these new type of filters we do not use this system of circles. Instead, we apply an iterative procedure. We start by applying a bank of DoG filters with different scales in the luminance channel. We find the position at which we obtain the global maximum response and use that position to form the first tuple $(\gamma_1, \sigma_1, \delta_1, \rho_1, \phi_1)$, where γ_1 is set to 1 that represents the luminance channel, σ_1 is the

scale of the DoG function that gives the maximum response at the considered location, δ_1 is the polarity (1 for center-on and -1 for center-off) of the concerned DoG function while ρ_1 and ϕ_1 are the distance and polar angle of that position with respect to the center of the prototype, respectively. Next, we remove the responses of all DoG functions in a circular region with a radius of $0.96\sigma_1$, which is the area responsible for by the selected DoG function. Then, we consider the next global maximum DoG response and form the second tuple. This procedure is repeated for all the three L , a and b channels until all the DoG responses are taken in consideration.

The application of a color-blob-based COSFIRE filter has exactly the same steps as edge-based COSFIRE filters, namely *filter*, *blur*, and *shift* for each tuple and *multiply* in the end. For detailed technical details we refer the reader to [26].

E. Forming descriptors

Suppose we have m edge-based COSFIRE filters and n color-blob-based COSFIRE filters, which are configured from randomly selected local patterns in training face images. The random regions selected as prototypes have a size of 19×19 pixels, in order to include significant shapes.

For a given image we apply the sets of m and n COSFIRE filters and for each COSFIRE filter response map we select features from a spatial pyramid with three levels. In the first level we consider the whole map as one tile and choose the maximum response across all of its locations. In the second level we split the response map into 2×2 tiles and take the maximum from each tile. Similarly, in the third layer, we consider a 4×4 grid of tiles and take the maximum response in every tile. In this way we end up with a $21m -$ and $21n -$ elements vectors for the edge-based and color-blob-based COSFIRE filters, respectively. Considering the edge-based COSFIRE filters, the first m elements are the maximum responses of the set of m edge-based filters, the second m elements are the maximum responses in the top left quadrant of the response maps of the same filters, and so forth. The same order is used to generate the color-blob-based descriptor.

The descriptor elements corresponding to the same tile are normalized to unit length. This means that the sum of squares of every descriptor is equal to 21.

F. Stacked classifier for gender recognition

We use a stacked classifier to fuse both type of descriptors. We first train two independent SVM classifiers with the below chi-squared kernel $K(x_i, y_j)$:

$$K(x_i, y_j) = \frac{(x_i - y_j)^2}{\frac{1}{2}(x_i + y_j) + \epsilon} \quad (1)$$

where x_i and y_j are the i -th and j -th training examples and ϵ is a very small number¹ used to avoid numerical errors. Then we apply these two SVM classifiers to the training images and form a descriptor that contains the scores of both classifiers. Finally, we use the new 2-element vectors to train a new SVM classifier with a linear kernel.

¹In Matlab we use the in-built function *eps*

TABLE I: Number of training and test images, with details about male and female faces, used for experiments on the GENDER-COLOR-FERET dataset.

Training set	Test set	Total	
209	209	418	M
209	209	418	F
418	418	836	Total

Fig. 3: Examples of GENDER-COLOR-FERET images.

III. EXPERIMENTAL RESULTS

A. Dataset

The dataset used for our experiments is the GENDER-COLOR-FERET, recently proposed in [16] and publicly available². Table I reports the composition of the dataset, which consists of 836 images, of which 418 are male faces and 418 are female faces. The dataset is divided equally between training and test set and the same person does not appear in both sets. Figure 3 shows some samples of the dataset. The face images are frontal with respect to the camera, but exhibit variations in terms of distance, expression, illumination, background, age and race.

B. Evaluation and results

Table II reports the results achieved by the proposed method on the GENDER-COLOR-FERET dataset. The edge-based COSFIRE method with 180 filters, configured as described in [24], is able to correctly recognize 95.2% of the face images in the test set, confirming its suitability for gender classification. In order to have the system balanced we configure 180 color-blob-based COSFIRE filters with the default parameters described in [26], except for λ which we set to 1.75. This value was determined from a grid search on the training set. Such algorithm achieves an accuracy of 94.7% on the test set, slightly lower than the performance obtained by the edge-based COSFIRE method. The complementarity of the two types of features is demonstrated by the results achieved with the stacked classifier. The combination between edge- and color-blob-based methods achieves an accuracy rate of 96.4%, significantly higher than the performance reached by the individual experts.

Table III shows the comparison of the results with other already published gender recognition methods. All the classifiers

TABLE II: Results of the edge-based and the color-blob-based COSFIRE methods on the GENDER-COLOR-FERET dataset.

Method	Accuracy (%)
Edge-based COSFIRE	95.2
Color-blob-based COSFIRE	94.7
Proposed	96.4

TABLE III: Comparison of the results with other scientific methods on the GENDER-COLOR-FERET dataset.

Method	Accuracy (%)
Magnitude descriptors [16]	83.0
RAW, LBP, HOG [15]	93.3
LBP, HOG [15]	95.2
Edge-based COSFIRE [24]	95.2
Proposed	96.4

have been trained and evaluated with the same procedure, so as to make a completely fair comparison. The classifier based on the local magnitude descriptors described in [16], suitable for real-time processing of video sequences, is the less effective on the GENDER-COLOR-FERET test images (83.0%). The classifier that takes input from handcrafted features proposed in [15] is able to achieve an accuracy of 93.3%, which is further improved by removing the RAW features (95.2%), that are sensitive to small image variations. The proposed method outperforms all the considered competitors, including the edge-based COSFIRE proposed in [24].

IV. DISCUSSION

The high recognition results that we achieve demonstrate the effectiveness of the proposed method for the problem at hand. It also outperforms existing methods in the considered dataset.

The edge-based COSFIRE filters confirm their suitability for discerning between male and female facial traits. So the classifier based on these trainable shape features is able to reproduce the cognitive process of the brain for the recognition of gender from face images, also in absence of obvious cues.

The color-blob-based COSFIRE filters, applied for the first time in the face analysis field, demonstrate their capability to recognize the gender by encoding information about the spatial arrangement of different shades of colors, which characterize various parts of the face, such as eyes, cheekbones and lips. These differences are even more evident in presence of make-up and allow the classifier trained with such features to successfully distinguish between men and women.

The experimental evaluation confirms our hypothesis that the two types of features are complementary to each other, since they are able to capture different aspects of the human face. The potential of the multi-expert system is further confirmed by another experiment that we did to evaluate the performance achievable by combining the decisions of the two experts. Such experiment proves that an ideal

²<http://mivia.unisa.it>

combination rule, namely the one which is able to take the correct decision when at least one of the classifiers is right, may achieve an accuracy of 98.3%, higher than the one obtained by the proposed method (96.4%). During our experiments we evaluated other fusion techniques, namely the majority and the weighted voting, and finally chose the stacked classification scheme as it achieved the best results. This evidence suggests that a future direction may be the investigation of a more effective combination rule, with the aim of achieving the ideal performance.

The main advantage of the proposed method is that it does not require the engineering of features that better recognize the gender. It uses trainable COSFIRE filters that are configured from randomly selected local patterns in training images. Moreover, the proposed method does not restrict the use of a specific number of COSFIRE filters. While in this work we used 180 edge-based and 180 color-blob-based COSFIRE filters, the method allows the use of any number of COSFIRE filters. The computational complexity does not increase linearly with the number of filters used. This is because there are many computations that are shared among tuples of different COSFIRE filters. COSFIRE filters are also highly parallelizable as the computations required by the tuples are independent of each other.

There are various future directions of the proposed approach. One direction would be to investigate the automatic selection of the most discriminative COSFIRE features. In this work, we rely on the random selection of local patterns from training images to configure COSFIRE filters. This procedure, however, does not guarantee the configuration of the most effective set of COSFIRE filters. So far COSFIRE filters have always been configured with single prototype patterns including this work. Another direction, therefore, would be to learn the tolerance behaviour of COSFIRE filters with gradient descent, for instance, on the training images rather than setting the same tolerance for all filters invariably. Thirdly, we are eager to extend this work by investigating other recognition tasks from face images, such as ethnicity, expression and age categories.

V. CONCLUSION

The stacked classification method that fuses edge-based and color-blob COSFIRE features is effective for the recognition of gender in face images and outperforms existing methods on the same data set. The edge-based features are robust in the encoding of shape characteristics of faces while the color-blob-based features are able to encode diffuse boundaries in different color channels which may be more pronounced in females. The two types of features turn out to be complementary in the recognition task at hand as the accuracy rate improves significantly when combined together.

The trainable nature of the COSFIRE features along with the automatic procedure that configures the filters and forms the face descriptor, allows the method to be easily adapted to

other face analysis problems, such as ethnicity recognition, age estimation and expression analysis.

REFERENCES

- [1] C. B. Ng, Y. H. Tay, and B.-M. Goi, *Recognizing Human Gender in Computer Vision: A Survey*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2012, pp. 335–346.
- [2] T. Ojala, M. Pietikäinen, and T. Mäenpää, “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.
- [3] H.-C. Lian and B.-L. Lu, “Multi-view gender classification using local binary patterns and support vector machines,” in *Advances in Neural Networks-ISNN 2006*. Springer, 2006, pp. 202–209.
- [4] Z. Yang and H. Ai, “Demographic classification with local binary patterns,” in *Advances in Biometrics*. Springer, 2007, pp. 464–473.
- [5] 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.
- [6] B. Moghaddam and M. Yang, “Learning gender with support faces,” *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 24, no. 5, pp. 707–711, 2002.
- [7] S. Baluja and H. A. Rowley, “Boosting sex identification performance,” *International Journal of computer vision*, vol. 71, no. 1, pp. 111–119, 2007.
- [8] J. Yang, D. Zhang, A. F. Frangi, and J. Y. Yang, “Two-dimensional pca: a new approach to appearance-based face representation and recognition,” *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 26, no. 1, pp. 131–137, 2004.
- [9] G. Rhodes, C. Hickford, and L. Jeffery, “Sex-typicality and attractiveness: Are supermale and superfemale faces super-attractive?” *British Journal of Psychology*, vol. 91, no. 1, pp. 125–140, 2000.
- [10] N. Dalal and B. Triggs, “Histograms of oriented gradients for human detection,” in *Computer Vision and Pattern Recognition, 2005. CVPR 2005. IEEE Computer Society Conference on*, vol. 1. IEEE, 2005, pp. 886–893.
- [11] V. Singh, V. Shokeen, and M. B. Singh, “Comparison of feature extraction algorithms for gender classification from face images,” in *International Journal of Engineering Research and Technology*, no. 5 (May-2013). ESRSA Publications, 2013.
- [12] G. Guo, C. R. Dyer, Y. Fu, and T. S. Huang, “Is gender recognition affected by age?” in *Computer Vision Workshops (ICCV Workshops), 2009 IEEE 12th International Conference on*. IEEE, 2009, pp. 2032–2039.
- [13] L. A. Alexandre, “Gender recognition: A multiscale decision fusion approach,” *Pattern Recognition Letters*, vol. 31, no. 11, pp. 1422–1427, 2010.
- [14] J. E. Tapia and C. A. Perez, “Gender classification based on fusion of different spatial scale features selected by mutual information from histogram of lbp, intensity, and shape,” *Information Forensics and Security, IEEE Transactions on*, vol. 8, no. 3, pp. 488–499, 2013.
- [15] G. Azzopardi, A. Greco, and M. Vento, “Gender recognition from face images using a fusion of svm classifiers,” in *Image Analysis and Recognition: 13th International Conference, ICIAR 2016, in Memory of Mohamed Kamel, Póvoa de Varzim, Portugal, July 13-15, 2016, Proceedings*. Springer International Publishing, 2016, pp. 533–538.
- [16] G. Azzopardi, A. Greco, A. Saggese, and M. Vento, “Fast gender recognition in videos using a novel descriptor based on the gradient magnitudes of facial landmarks,” in *Advanced Video and Signal Based Surveillance (AVSS), 2017 14th IEEE International Conference on*. IEEE, 2017, pp. 1–6.
- [17] Y. S. El-Din, M. N. Moustafa, and H. Mahdi, “Landmarks-sift face representation for gender classification,” in *International Conference on Image Analysis and Processing*. Springer, 2013, pp. 329–338.
- [18] N. Wang, X. Gao, D. Tao, and X. Li, “Facial feature point detection: A comprehensive survey,” *arXiv preprint arXiv:1410.1037*, 2014.
- [19] G. Levi and T. Hassner, “Age and gender classification using convolutional neural networks,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2015, pp. 34–42.
- [20] J. van de Wolfshaar, M. F. Karaaba, and M. A. Wiering, “Deep convolutional neural networks and support vector machines for gender recognition,” in *Computational Intelligence, 2015 IEEE Symposium Series on*. IEEE, 2015, pp. 188–195.

- [21] R. Ranjan, V. M. Patel, and R. Chellappa, "Hyperface: A deep multi-task learning framework for face detection, landmark localization, pose estimation, and gender recognition," *arXiv preprint arXiv:1603.01249*, 2016.
- [22] S. Jia and N. Cristianini, "Learning to classify gender from four million images," *Pattern Recognition Letters*, vol. 58, pp. 35–41, 2015.
- [23] G. Azzopardi and N. Petkov, "Trainable COSFIRE filters for keypoint detection and pattern recognition," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 35, no. 2, pp. 490–503, 2013.
- [24] G. Azzopardi, A. Greco, and M. Vento, "Gender recognition from face images with trainable cosfire filters," in *Advanced Video and Signal Based Surveillance (AVSS), 2016 13th IEEE International Conference on*. IEEE, 2016, pp. 235–241.
- [25] G. Azzopardi, A. Greco, A. Saggese, and M. Vento, "Fusion of domain-specific and trainable features for gender recognition from face images," *IEEE Access*, 2018.
- [26] B. Gecer, G. Azzopardi, and N. Petkov, "Color-blob-based cosfire filters for object recognition," *Image and Vision Computing*, vol. 57, pp. 165–174, 2017.
- [27] P. Viola and M. J. Jones, "Robust real-time face detection," *International journal of computer vision*, vol. 57, no. 2, pp. 137–154, 2004.