Towards Face Unlock: On the Difficulty of Reliably Detecting Faces on Mobile Phones

Rainhard D. Findling
Department for Mobile Computing
Upper Austria University of Applied Sciences
Softwarepark 11
Hagenberg, Austria
rainhard.findling@fh-hagenberg.at

Rene Mayrhofer
Department for Mobile Computing
Upper Austria University of Applied Sciences
Softwarepark 11
Hagenberg, Austria
rene.mayrhofer@fh-hagenberg.at

ABSTRACT

Currently, reliable face detection and recognition are becoming more important on mobile devices – e.g. to unlock the screen. However, using only frontal face images for authentication purposes can no longer be considered secure under the assumption of easy availability of frontal snapshots of the respective device owners from social networks or other media. In most current implementations, a sufficiently high-resolution face image displayed on another mobile device will be enough to circumvent security measures. In this paper, we analyze current methods to face detection and recognition regarding their usability in the mobile domain, and then propose an approach to a Face Unlock system on a smart phone intended to be more secure than current approaches while still being convenient to use: we use both frontal and profile face information available during a pan shot around the user's head, by combining camera images and movement sensor data. Current results to face detection are promising, but reliable face recognition needs further research.

Categories and Subject Descriptors

I.5 [**Pattern Recognition**]: Clustering, Applications, Implementation

Keywords

Face detection; face recognition; mobile phone; user authentication

1. INTRODUCTION

Face detection and face recognition can support a wide variety of use cases: Face detection – finding the exact (or estimated) position of human faces in pictures – is, among many others, used in current digital cameras for adjusting depth of focus, in video surveillance systems, or as basis for face recognition. Face recognition – identifying people by using information from their faces – is commonly used for video surveillance, but is now becoming a highly attractive alternative for authenticating users to their own mobile devices.

In contrast to the current standard approach of PIN, password, or swipe gesture entry to unlock mobile devices, biometric methods such as face recognition are immune to the shoulder surfing attack (an attacker watching the display and therefore the textual/visual passcode entry while the user authenticates). Unfortunately, using only frontal face information for face-based authentication is still easy to circumvent e.g. by presenting a sufficiently large and high-quality photograph of the person to the camera [1, 4, 19, 22]. Therefore, our aim is to combine all face information that is available from moving the mobile device 180° from left profile over frontal to right profile of the user's face. This approach is still fast and convenient to use but harder to attack, as more information than contained in a frontal picture of the face would be needed (i.e., attackers would need to provide a 3D reconstruction of the person's face or a closely synchronized video stream instead of a single, static photograph that can e.g. often be found on social network websites).

In this paper, we make the following novel contributions:

- First, we analyse current approaches to face detection and discuss their issues especially with regards to their use on hand-held mobile devices for self-authentication (section 2).
- Then we present an approach to detecting faces from a non-frontal camera viewpoint within the difficult environment of hand-held smart phone cameras (section 3).
- Finally, we show in a practical implementation under Android that our approach supports sufficiently reliable face detection for a roughly 180° range (section 3.3) as a prerequisite for future work on pseudo-3D face recognition for more secure implementations of Face Unlock.

Although we briefly describe our prototype implementation of the face recognition step (section 3.4) to provide context for the specifics of our approach to face detection, this is not the main focus of the present paper and only applies current state-of-the-art classification methods. Improvements to the face recognition step and a move from multiple 2D images to a (pseudo-) 3D space is subject to future work.

We note that, as with every approach to using biometric information for authentication, no key revoke is possible (i.e.

MoMM2012, 3-5 December, 2012, Bali, Indonesia

when the stored template or reference images were compromised and the authentication data would therefore need to be changed). Our approach therefore does not target high security systems, but is intended to be used for personal devices that are in frequent use and where this approach still provides a higher security level than current approaches and is more convenient to use.

2. ANALYSIS OF CURRENT APPROACHES

2.1 Face Detection

Turn and Pentland [20] initially proposed the use of principal component analysis (PCA) and Eigenface subspace for face detection and recognition. Sung and Poggio [18] used view-based model clusters that distinguish between "face" and "non-face". Rowley et. al. [13] use neural networks to estimate if parts of an image contains a face. Schneiderman and Kanade [17] identify objects – including faces – by using wavelet transformation. Sahoolizadeh et. al. [14] combines Gabor wavelets and neural networks for face detection and recognition. Bayesian discriminating features were used by Liu [10], which compares likelihood density estimations of an image to decide if an image contains a face. Finally, the use of skin color for face detection was investigated by different authors, e.g. [6, 11, 26] but turned out to be less reliable than other approaches.

Viola and Jones [21] proposed cascades of boosted classifiers, which is now considered one of the standard approaches to face detection. With the extension by Lienhart and Maydt [9], this approach to face detection is used by default in OpenCV. The boosted classifiers are trained for effective, haarlike features from positive and negative face examples, which have to be provided during a training phase. Those features represent oriented contrast in an image to identify faces in the next step. The classifiers constructed are "weak" classifiers as the basis for standard boosting to obtain "strong" classifiers. Those strong classifiers are then combined to form a classifier cascade. The faster classifiers of the cascade are applied first in the processing chain: as soon as an image fails to pass one classifier, it will not be presented to subsequent classifiers. Finally, to find faces of different sizes, a sliding window principle is applied. We are currently using the implementation of this algorithm in OpenCV for our approach.

However, many of the existing face detection and recognition mechanisms were designed to work for frontal face images, without considering profile face images in the first place. One approach to use those mechanisms for more than frontal face images was mentioned by Pentland [12], who used different models for different points of view.

2.2 Face Recognition

Many different approaches to face recognition have been published in the past 20 years, and for a more comprehensive review we refer to [7, 23, 27, 28]. One of the most important was using Eigenfaces for face detection and recognition by Turk and Pentland [20], which is still used as a baseline for benchmarking newer approaches to face recognition. As the established standard (baseline) approach, we also apply it in our first experiments. Eigenfaces are more prone to recognition errors caused by changes in illumination than Fisherfaces, as shown by Belhumeur et. al. [2]. Other approaches than dimensionality reduction were also proposed, like Deformable Templates by Yuille [25] or Elastic Bunch Graph Mapping by Wiskott et. at. [24].

The recent approaches of Sahoolizadeh et. al. [14] and Kurutach et. al. [8] seem interesting, as they both report a significantly improved reliability compared to the baseline Eigenfaces approach. Kurutach et. al. use Trace Transform to obtain features from faces, then the Hausdorff Distance and Shape Context are used to compare faces. Their system has been trained and tested with photos from the ORL [15] and YALE [5] data base, where photos used for tests have been rotated and scaled arbitrarily before usage. The measured detection rate is reported to be better than when using Eigenfaces in all cases, with the worst detection rate being 68.87% for photos containing a scream face expression. Unfortunately, this approach cannot be applied to profile faces, as detected faces get normalized in scale and rotation based on the eyes' positions. Sahoolizadeh et. al. [14] combine Gabor wavelets and neural networks for reliable face detection and recognition. Their system has been trained and tested with photos from the ORL data base, with a reported detection rate of 77.9% to 90.3%. Sahoolizadeh et. al. state that the current system only detects upright faces looking at the camera – and that future work will include separate versions of the system for detection and recognition at different head positions.

The techniques described in the present paper are applicable to all approaches to face recognition, as most of them require robust face detection as a first step. Therefore, by improving upon previous work for non-frontal face detection, we provide groundwork for future improvements towards non-frontal or mixed face recognition (as discussed below).

3. IMPROVING FACE RECOGNITION ON MOBILE PHONES

3.1 Intended Usage

The Face Unlock we intend to develop requires a mobile device with a frontal camera and sensors such as a gyroscope. Although our mid-term aim in terms of usability is a pseudo-3D reconstruction of facial features with a quick, non-standardized swipe of the user's mobile phone around the front side of her/his head, in our current state of prototypical implementation we require the user to perform a more formalized swipe of the camera: the user holds her/his mobile phone either right or left of her/his head, so that the frontal camera points towards the ear. The arm holding the phone should be stretched. The user then moves the mobile phone in a rough circle via the frontal view along to the other side of her/his head, so that the frontal camera points towards the other ear. The arm holding the phone should be kept stretched. The data obtained by the mobile phone, including a frontal camera video stream and motion sensor time series, is then used for Face Unlock to avoid the simple attack vector of presenting a static picture of the user's face to a static phone.

3.2 Environment

As environment for a Face Unlock prototype we are targeting an state of the art mobile phone with frontal camera and at least a gyroscope and based on Android to enable future integration into the platform's unlock feature.

For our current implementation, we use a Google/Samsung Nexus S GT-I9023 device, currently running Android 2.3.3. For face detection and recognition we use OpenCV compiled for Android[1] and JavaCV for Android[2] as wrapper around OpenCV.

3.3 Face Detection

The Face Unlock application has a state `STATE`, which initially is `IDLE`. As the user holds the mobile phone with the frontal camera towards one ear, the application changes from `IDLE` to `ACTIVE`. The application stays active as the user moves the mobile phone via his frontal face towards the other ear. As soon as the frontal camera points to the other ear – determined by the gyroscope data –, the application goes from `ACTIVE` to `IDLE` again. In our current implementation, changing `STATE` is done by the user pressing a button. As long as the application is `ACTIVE`, photos are taken using the frontal camera. The application decides when the next photo should be taken by monitoring the device angle, resulting from the gyroscope time series. If the changes in the device angle since the last photo are larger than a defined threshold α, the next photo is taken. For our experiments, $\alpha = 15°$ has been used. Each photo is stored along with metadata (most importantly the current device angle). Therefore, roughly the same number of photos are made for a pan shot done for each Face Unlock, and processing the photos can be done afterwards.

We do not currently record a full video stream of the whole camera movement across the user's face because of the mentioned limitations in the mobile phone APIs: on the one hand, most phones offer only limited resolution in video mode when compared to picture mode, and on the other hand, Android does not yet support accessing the raw video stream with low processing overhead from third-party applications. Additionally, the limited processing resources on current mobile phones would not allow to process the full video stream for face recognition in real time.

As the application switches from `ACTIVE` to `IDLE`, the following steps are processed: first, a normalization of the metadata stored with each photo is performed. Assuming that – seen from a frontal face perspective – the user has held the mobile phone at roughly the same angle when starting and ending the Face Unlock, the frontal face perspective is defined to be at an angle of 0°. When β is the total angle the mobile phone has rotated, the normalization is performed so that the maximum left angle of all photos is roughly $-\frac{\beta}{2}$, and the maximum right angle of all photos is roughly $\frac{\beta}{2}$. Second, all photos are converted to gray scale. This conversion incurs some information loss, but most face recognition algorithms operate on gray scale only to be more robust against different lighting conditions, and the limitation to a single channel allows faster processing in subsequent stages.

Finally, face detection is performed for each photo. The OpenCV face detection classifier cascades is chosen depending on the metadata stored along with each photo, where γ is the device angle the photo was shot at and ϕ is a predefined threshold angle. If $\gamma < -\phi$, the `PROFILE` classifier cascade gets chosen. If $\gamma > \phi$, the picture is mirrored[3] and the `PROFILE` classifier cascade is chosen. If $|\gamma| \leq \phi$, the `FRONTAL-ALT` classifier cascade is chosen. For our experiments $\gamma = 30°$ was used. Face detection is then performed using the chosen classifier. Finally, areas that are found to contain a face are extracted from the pictures and saved to separate face images along with the angle the picture has been taken at. Figure 1 shows the pictures recorded during one pan shot, along with the faces detected in those pictures. These face images are then used for face recognition in the next step.

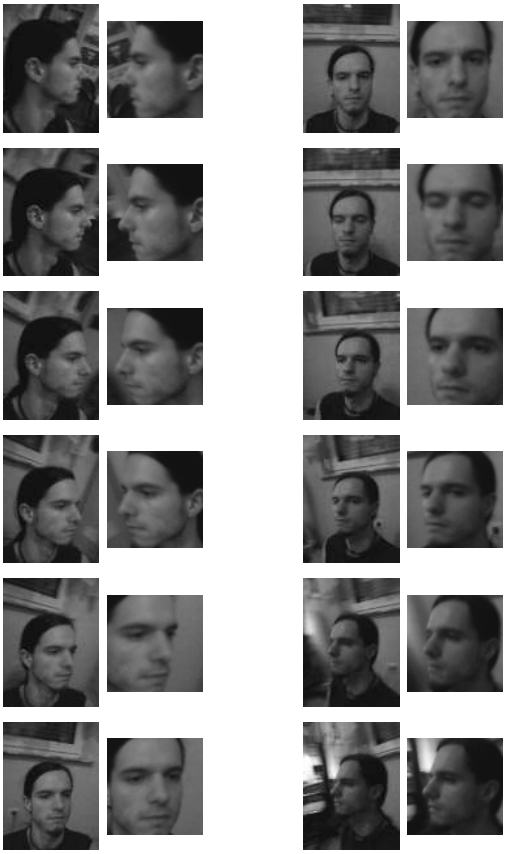

Figure 1: Pictures recorded and faces detected from one pan shot.

3.4 Face Recognition

For face recognition, the Face Unlock application currently contains several classifiers. Each classifier covers a certain angle-of-view α of the user's face, which corresponds to the multi-view approach of Pentland et. al. [12]. Therefore, face images shot at a similar angles will be assigned the same or a neighboring classifier in the normal case. For our experiments, we used $\alpha = 20°$, which results in about 9 classifiers for an assumed total device rotation of 180°.

The Face Unlock application can either be in `TRAIN` or in `CLASSIFY` mode. For both modes, the application takes the face images resulting from section 3.3. In `TRAIN`, the classifiers are trained with face images of people that should later be recognized. Therefore the identity of the person is set manually in this mode. Detected face images are assigned to a classifier, corresponding to their angle. This is required for the currently used Eigenfaces approach as pro-

[1] `http://opencv.alekcac.webfactional.com/downloads.html`

[2] `http://code.google.com/p/javacv/downloads/list`

[3] The OpenCV `PROFILE` classifier cascade only detects left profile faces.

posed by Turk and Pentland [20]. Eigenfaces are based on an average face, which has to be recalculated every time the training faces are expanded — otherwise, no expansion will be possible. As the user switches the Face Unlock application from `TRAIN` to `CLASSIFY`, each classifier is trained with all face images assigned to it. This training is done on the end-user phone without requiring server-assisted ("cloud") computation for privacy reasons. Our current implementation is therefore also a proof of concept of the feasibility of on-device biometric authentication on current smart phones.

When the Face Unlock application is in `CLASSIFY` mode, detected face images are classified by the classifier corresponding to the angle at which the face image has been shot. For each face image to classify, a classifier delivers a list of distances. Each distance corresponds to the difference of the face image to classify to the face images of the people known to the classifier. Probabilities of how certain the person currently unlocking the device is a person known to the system can be derived from these distance lists.

For our current, proof-of-concept implementation, we are summing up the probabilities of different angles to obtain an overall probability, with which access to the mobile phone can then either be granted or denied.

3.5 Face Database

For doing comprehensive tests on our face recognition approach, we have created a preliminary face database at FH Hagenberg. This database features 38 people with 1-3 pan shot image sets each and 95 such sets in total. Each set contains 4-5 images, shot at the angles $90°$, $45°$, $0°$, $-45°$ and $-90°$, with $0°$ being the frontal face perspective. The facial expressions of all images in a set are either normal, eyes closed or smiling, and the illumination of the faces is evenly good. One such set is shown in figure 2.

Figure 2: Pan shot image set from our preliminary face database.

For research purposes our preliminary face database is available on request via e-mail. An extended database with additional images and more sensor data for pan-shot movies will be made available at a later stage.

There are several reasons why we created this face database for testing our Face Recognition: a) the illumination of the faces shown in pictures taken with a pan shot around the users head vary strongly with each pan shot. Therefore, the test results would not be reproducible and comprehensive enough. b) in our experiments, the frontal camera photo quality strongly depended on how fast the user moved the mobile phone. Moving the mobile phone from one ear, along the frontal face perspective to the other ear took about 4 seconds to obtain photos of good quality. In case the user moved the mobile phone faster, the image quality was lowered due to motion blur, which consistently lowered the system's reliability. c) we are not aware of any other face databases available for research that contain face pan shots, state the angle at which a picture was taken, and have multiple pictures per angle and person available at the same time.

3.6 Current Results

Using the images from the our preliminary face database as input to our Face Recognition system results in a face detection rate (true positives in terms of authentication systems) of 100% for frontal face images (which use the `FRONTAL -ALT` classifier) and 90.5% for face images shot at angle γ and $|\gamma| = 45°$ and $|\gamma| = 90°$ (which use the `PROFILE` classifier). A few false positive cases from the `PROFILE` face detection, such as the examples shown in figure 3, negatively influence the latter face recognition, as they get used for training and test data as if they were correct results.

Figure 3: Examples for false positive results from the `PROFILE` face detection.

In the regular case a face can be detected in each picture taken in the pan shot, assuming a slow enough device movement of about 4 seconds for the total pan shot, and a sufficient illumination of the face, as shown in figure 1. In case of poorer illumination of the face, for some recorded pictures no faces might be detected, as shown in figure 4.

Even if this detection rate is sufficient for our current usage, improvements to the face detection might become necessary in the future, as more intensive tests of the algorithm in [3, 16]have shown that specially the profile face detection classifier of the OpenCV implementation suffers from a decreased detection rate.

The face recognition rates of our Face Unlock application have been evaluated using our preliminary face database in a test classification as follows:

1. Randomly chose a pan shot as test set. The person shown in this set is the test subject.
2. Chose other pan shots of the test subject as training sets.
3. Further add random pan shots of other persons to the training sets, until the training set contains 20 pan shots.
4. Train the Face Unlock application with the training set. The classification problem gets reduced to a binary classification problem by treating all images that belong to the test subject as being part of the positive class, and all images of all other persons as being part of a single negative class.
5. Test the Face Unlock application with the test set.

For 100 such classifications, the test subject got recognized in 78.5% using frontal face information only, and in 55.8% using pan shot face information. We argue that the overall recognition rate (i.e. the true positives rate for the

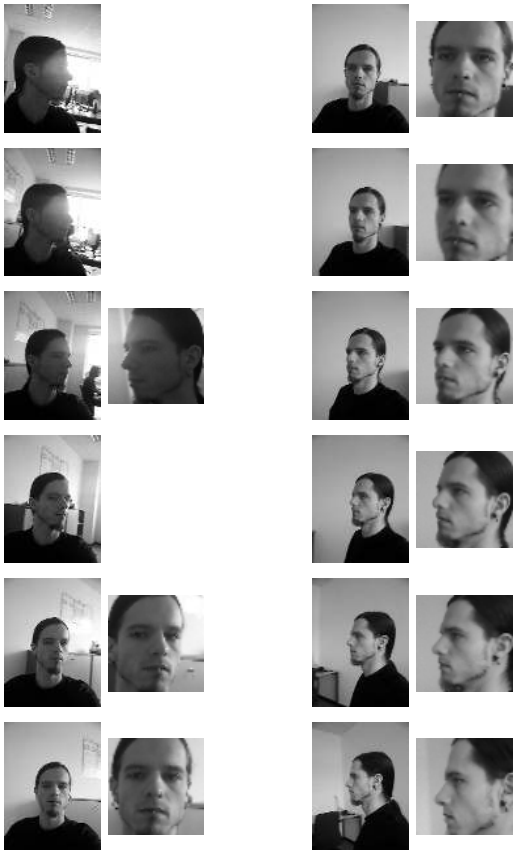

Figure 4: Pictures recorded and faces detected from one pan shot with more difficult illumination conditions.

authentication case) is lower when using pan shot information because of our current combination of probabilities of the different views by simply summing up: as frontal face pictures seem to be easier separable than profile face pictures, as stated e.g. by Santana et. al. [16], and the profile faces are being detected less reliable at the same time, the better results of frontal face recognition are extended by a 4 times larger amount of worse results from profile face recognition. However, we assume a significantly higher difficulty level for tricking the system into authentication when relying on the pan shots instead of only frontal face shots. An attacker would have to replay a synchronized video stream while moving the attacked device or manufacture a 3D bust of the owner's face. Although we can not yet quantify the resulting increase in security, we argue that a small decrease in recognition rate is outweighed by the increase in security, which would support the day-to-day use of Face Unlock even for application scenarios with higher security demands.

Our current results strongly indicate that face detection can be done sufficiently reliable even for pan shots, but that the second step of recognition based on Eigenfaces does not (yet) work reliably enough for further usage. Additionally, as stated by Belhumeur [2], changes in illumination are a major problem for recognition based on Eigenfaces. As changes in illumination are omnipresent in the mobile domain, this approach is not sufficient. As mentioned before, we implemented the Eigenfaces approach as a baseline and first step and are currently working on porting other algorithms to Android for further improvements to the face recognition step.

4. CONCLUSION

We are working on a Face Unlock application that aims to use most information available to a mobile phone from a pan shot around the user's head. The system is intended to be convenient to use, reliable and harder to attack than using frontal face information only. For processing face information from different points of view, we use multiple classifiers for face detection and recognition. The disadvantage is that we require more training data. However, we argue that the significant gain in security outweighs the longer training phase, because training only needs to be performed once for each user. Metadata, obtained from sensors such as a gyroscope, is used to determine which classifier to use for which information.

Currently, we are using the Viola and Jones algorithm implemented in OpenCV for face detection with cascades optimized for frontal and for side images, and the Eigenface approach of Turk and Pentland for face recognition. While the face detection seems to be reliable enough for further usage, the face recognition is error prone mainly due to changing illumination – which is ubiquitous in the mobile domain and one of the main problems identified by our work. Future work will have to explicitly deal with the issues of foreground and background illumination, changing background, and below-the-phase camera positions due to the typical position of mobile phones when held by end-users.

At the moment, we are extending the system to use different face recognition approaches. Further, we intend the usage of 3D information obtained from stereo cameras, for both face detection and recognition.

Acknowledgement

This work is supported by the Austrian national funding line FFG COIN within project SESAME-S. The images contained in our preliminary face database have been taken by Christine Aigner in a high quality and colored form as part of her bachelor thesis and got preprocessed to the form stated in section 3.5 by Rainhard Findling.

5. REFERENCES

[1] W. Bao, H. Li, N. Li, and W. Jiang. A liveness detection method for face recognition based on optical flow field. In *Image Analysis and Signal Processing, 2009. IASP 2009. International Conference on*, pages 233 –236, Apr. 2009.

[2] P. N. Belhumeur, J. a. P. Hespanha, and D. J. Kriegman. Eigenfaces vs. fisherfaces: Recognition using class specific linear projection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19(7):711–720, July 1997.

[3] D. Douxchamps and N. Campbell. Robust real time face tracking for the analysis of human behaviour. In A. Popescu-Belis, S. Renals, and H. Bourlard, editors, *Machine Learning for Multimodal Interaction*, volume 4892 of *Lecture Notes in Computer Science*, pages 1–10. Springer Berlin / Heidelberg, 2008.

[4] R. Frischholz and A. Werner. Avoiding replay-attacks in a face recognition system using head-pose estimation. In *Analysis and Modeling of Faces and Gestures, 2003. AMFG 2003. IEEE International Workshop on*, pages 234 – 235, Oct. 2003.

[5] A. Georghiades, P. Belhumeur, and D. Kriegman. From few to many: Illumination cone models for face recognition under variable lighting and pose. *IEEE Trans. Pattern Anal. Mach. Intelligence*, 23(6):643–660, 2001.

[6] R. Hsu, M. Abdel-Mottaleb, and A. K. Jain. Face detection in color images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24:696–706, 2002.

[7] A. K. Jain and S. Z. Li. *Handbook of Face Recognition*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 2005.

[8] W. Kurutach, R. Fooprateepsiri, and S. Phoomvuthisarn. A highly robust approach face recognition using hausdorff-trace transformation. In *ICONIP (2)*, pages 549–556, 2010.

[9] R. Lienhart and J. Maydt. An extended set of haar-like features for rapid object detection. In *IEEE ICIP 2002*, pages 900–903, 2002.

[10] C. Liu. A bayesian discriminating features method for face detection. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 25(6):725 – 740, june 2003.

[11] B. Martinkauppi. *Face colour under varying illumination - analysis and applications.* PhD thesis, University of Oulu, 2002.

[12] A. Pentland, B. Moghaddam, and T. Starner. View-based and modular eigenspaces for face recognition. In *Computer Vision and Pattern Recognition, 1994. Proceedings CVPR '94., 1994 IEEE Computer Society Conference on*, pages 84–91, June 1994.

[13] H. A. Rowley, S. Member, S. Baluja, and T. Kanade. Neural network-based face detection. *IEEE Transactions On Pattern Analysis and Machine intelligence*, 20:23–38, 1998.

[14] H. Sahoolizadeh, D. Sarikhanimoghadam, and H. Dehghani. Face detection using gabor wavelets and neural networks. World Academy of Science, Engineering and Technology 45, 2008.

[15] F. Samaria and A. Harter. Parameterisation of a stochastic model for human face identification. In *Applications of Computer Vision, 1994., Proceedings of the Second IEEE Workshop on*, pages 138 –142, Dec. 1994.

[16] M. C. Santana, O. Déniz-Suárez, L. Antón-Canalís, and J. Lorenzo-Navarro. Face and facial feature detection evaluation - performance evaluation of public domain haar detectors for face and facial feature detection. In A. Ranchordas and H. AraÃžjo, editors, *VISAPP (2)*, pages 167–172. INSTICC - Institute for Systems and Technologies of Information, Control and Communication, 2008.

[17] H. Schneiderman and T. Kanade. Object detection using the statistics of parts. *International Journal of Computer Vision*, 56(3):151–177, Feb. 2004.

[18] K. Sung and T. Poggio. Example-based learning for view-based human face detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20:39–51, 1998.

[19] R. Tronci, D. Muntoni, G. Fadda, M. Pili, N. Sirena, G. Murgia, M. Ristori, and F. Roli. Fusion of multiple clues for photo-attack detection in face recognition systems. In *Biometrics (IJCB), 2011 International Joint Conference on*, pages 1–6, Oct. 2011.

[20] M. Turk and A. Pentland. Eigenfaces for recognition. *Cognitive Neuroscience*, 3(1):71–86, Jan. 1991.

[21] P. Viola and M. Jones. Rapid object detection using a boosted cascade of simple features. *Proceedings CVPR*, 1:511–518, 2001.

[22] M. Wagner and G. Chetty. Liveness assurance in face authentication. In S. Z. Li and A. Jain, editors, *Encyclopedia of Biometrics*, pages 908–916. Springer US, 2009.

[23] H. Wechsler. *Reliable Face Recognition Methods: System Design, Implementation and Evaluation (International Series on Biometrics)*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 2006.

[24] L. Wiskott, J.-M. Fellous, N. Krüger, and C. V. D. Malsburg. Face recognition by elastic bunch graph matching. *IEEE transactions on pattern analysis and machine intelligence*, 19:775–779, 1997.

[25] A. L. Yuille, P. W. Hallinan, and D. S. Cohen. Feature extraction from faces using deformable templates. *International Journal of Computer Vision*, 8(2):99–111, Aug. 1992.

[26] B. D. Zarit, B. J. Super, and F. K. H. Quek. Comparison of five color models in skin pixel classification. In *International Workshop on Recognition, Analysis, and Tracking of Faces and Gestures in Real-Time Systems*, pages 58–63, 1999.

[27] W. Zhao, R. Chellappa, P. J. Phillips, and A. Rosenfeld. Face recognition: A literature survey. *ACM Comput. Surv.*, 35(4):399–458, Dec. 2003.

[28] X. Zou, J. Kittler, and K. Messer. Illumination invariant face recognition: A survey. In *Biometrics: Theory, Applications, and Systems, 2007. BTAS 2007. First IEEE International Conference on*, pages 1 –8, Sept. 2007.