

A Realistic Simulation Tool for Testing Face Recognition Systems under Real-World Conditions*

M. Correa, J. Ruiz-del-Solar, S. Parra-Tsunekawa, R. Verschae
Department of Electrical Engineering, Universidad de Chile
Advanced Mining Technology Center, Universidad de Chile

Abstract. In this article, a tool for testing face recognition systems under uncontrolled conditions is proposed. The key elements of this tool are a simulator and real face and background images taken under real-world conditions with different acquisition angles. Inside the simulated environment, an observing agent, the one with the ability to recognize faces, can navigate and observe the real face images, at different distances, angles and with indoor or outdoor illumination. During the face recognition process, the agent can actively change its viewpoint and relative distance to the faces in order to improve the recognition results. The simulation tool provides all functionalities to the agent (navigation, positioning, face's image composing under different angles, etc.), except the ones related with the recognition of faces. This tool could be of high interest for HRI applications related with the visual recognition of humans, as the ones included in the RoboCup @Home league. It allows comparing and quantifying the face recognition capabilities of service robots under exactly equal working conditions. It could be a complement to existing tests in the RoboCup @Home league. The applicability of the proposed tool is validated in the comparison of three state of the art face recognition methods.

Keywords: Face Recognition, Face Recognition Benchmarks, Evaluation Methodologies, RoboCup @Home.

1 Introduction

Face recognition in controlled environments is a relative mature application field (see recent surveys in [1][2][3][4]). However, face recognition in uncontrolled environments is still an open problem [9][10]. Recent journal's special issues [6], workshops [7], and databases [8] are devoted to this topic. Main factors that still disturb largely the face recognition process in uncontrolled environments are [10][11]: (i) variable illumination conditions, especially outdoor illumination, (ii) out-of-plane pose variations, and (iii) facial expression variations. The use of more complex sensors (thermal-, high-resolution-, and 3D- cameras), 3D face models, illumination models, and sets of images of each person that cover various face variations are some of the approaches being used to deal with the mentioned drawbacks, in different application domains [10][11].

* This research was partially funded by FONDECYT under Project Number 1090250.

A very important component in the development of face recognition methodologies is the availability of suitable databases, benchmarks, and evaluation methodologies. For instance, the very well known FERET database [13], one of the most employed face databases that also includes a testing protocol, has been very important in the development of face recognition algorithms for controlled environments in the last years. Some relative new databases such as the LFW (Labeled Faces in the Wild) [8] and FRGC (Face Recognition Grand Challenge) [14][12], among others, intend to provide real-world testing conditions. In applications such as HRI (Human Robot Interaction) and surveillance the use of spatiotemporal context or active vision mechanisms in the face recognition process¹ can increase largely the performance of the systems. However, face recognition approaches that include these dynamic mechanisms cannot be validated properly using current face databases (see database examples in [5]). Even the use of video face databases does not allow testing the use of those ideas. For instance, in a recorded video it is not possible to change actively the observer's viewpoint. The use of a simulator could allow accomplishing this (viewpoint changes), however, a simulator is not able to generate faces and backgrounds that looks real/natural enough.

Nevertheless, the combined used of a simulation tool with real face and background images taken under real-world conditions, could allow to accomplish the goal of providing a tool for testing face recognition systems under uncontrolled conditions. In this case, more than providing a database and a testing procedure, the idea would be to supply a testing environment that provides a face database, dynamic image's acquisition conditions, active vision mechanisms, and an evaluation methodology. The main goal of this paper is to provide such a testing tool. This tool provides a simulated environment with persons located at different positions and orientations. The face images are previously acquired under different pitch and yaw angles², in indoor and outdoor variable lighting conditions. Inside this environment, an observing agent, the one with the ability to recognize faces, can navigate and observe the real face images (with real background information), at different distances, angles (yaw, pitch, and roll) and with indoor or outdoor illumination. During the recognition process the agent can actively change its viewpoint to improve the face recognition results. The simulation tool provides all functionalities to the agent, except the ones related with the recognition of the faces.

This testing tool could be of high interest for HRI applications related with the visual recognition of humans, as the ones included in the RoboCup @Home league. It allows comparing and quantifying the face recognition capabilities of service robots under exactly equal working conditions. In fact, the use of this testing tool could complement some of the real tests that are in use in the RoboCup @Home league.

This article is organized as follows. In section 2, related work in face databases and evaluation methodologies is outlined. In section 3, the proposed testing tool is described. Results of the applicability of the testing tool in the comparison of three state of the art face recognition methods are presented in section 4. Finally, some conclusions and projections of this work are presented in section 5.

¹ In this work we consider the face recognition process as the one composed by the face detection, face alignment, and face recognition stages.

² In-plane rotations can be generated by software (simulator).

2 Related Work

The availability of standard databases, benchmarks, and evaluation methodologies is crucial for the appropriate comparison of algorithms. There is a large amount of face databases and associated evaluation methodologies that consider different number of persons, camera sensors, and image acquisition conditions, and that are suited to test different aspects of the face recognition problem such as illumination invariance, aging, expression invariance, etc. Basic information about face databases can be found in [5][15].

The FERET database [13] and its associated evaluation methodology is a standard choice for evaluating face recognition algorithms under controlled conditions. Other popular databases used with the same purpose are Yale Face Database [16] and BioID [17]. Other database such the AR Face Database [18] and the University of Notre Dame Biometrics Database [19] include faces with different facial expressions, illumination conditions, and occlusions. From our point of view, all of them are far from considering real-world conditions.

The Yale Face Database B [20] and PIE [21] are the most utilized databases to test the performance of algorithms under variable illumination conditions. Yale Face contains 5,760 single light source images of 10 subjects, each seen under 576 viewing conditions (9 poses x 64 illumination conditions). For every subject in a particular pose, an image with ambient (background) illumination was also captured. PIE is a database containing 41,368 images of 68 people, each person under 13 different poses, 43 different illumination conditions, and with 4 different expressions. Both databases consider only indoor illumination.

The LFW database [8] consists of 13,233 images faces of 5,749 different persons, obtained from news images by means of a face detector. There are no eyes/fiducial point annotations; the faces were just aligned using the output of the face detector. The images of the LFW database have a very large degree of variability in the face's pose, expression, age, race, and background. However, due to LFW images are obtained from news, which in general are taken by professional photographers, they are obtained under good illumination conditions, and mostly in indoors.

FRGC ver2.0 database [12] consists of 50,000 face images divided into training and validation partitions. The validation partition consists of data from 4,003 subject sessions. A subject session consists of controlled and uncontrolled images. The uncontrolled images were taken in varying illumination conditions in indoors and outdoors. Each set of uncontrolled images contains two expressions, smiling and neutral.

3 Proposed Testing Tool

The proposed testing tool allows that an observing agent can navigate inside a virtual scenario, and observe a set of N persons. The faces of each of these persons are previously scanned under different yaw and pitch angles, and under different indoor and outdoor illumination conditions. This allows that every time that the agent observes a person's face at a given distance and viewpoint, the corresponding

images/observations are composed using a database of real faces and background images, instead of being generated by the simulator.

Considering that the goal of this tool is to test the recognition abilities of the agent, and not the navigation ones, navigation is simplified: the agent is placed in front of each person by the system. After being positioned, the agent analyzes its input images in order to detect and recognize human faces. Depending on the results of this analysis, the agent can change its relative pose. Every time that the agent changes its pose, the simulator composes the corresponding input images. The process is repeated until the agent observes all persons.

3.1 Image Acquisition System

Real face images are acquired at different yaw and pitch angles using a CCD camera mounted in a rotating structure (see diagram in fig. 1a). The person under scan is in a still position, while the camera, placed at the same height than the person's face and at a fixed distance of 140 cm, rotates in the axial plane (the camera height is adjustable). An encoder placed in the rotation axis calculates the face's yaw angle. There are not restrictions on the person's face expression. The system is able to acquire images with a resolution of 1° . However, in this first version of the testing tool, images are taken every 2° . The scanning process takes 25 seconds, and we use a 1280 x 960 pixels CCD camera (DFK 41BU02 model). In the frontal image, the face's size is about 200x250 pixels.

Variations in pitch are obtained by repeating the described processes, with the different required pitch angles are required. In each case, the camera height is maintained, but the person looks at a different reference points in the vertical axis, which are located at 160 cm in front of the person (see fig. 1a). In our experience, pitch angles of -15° , 0° , and 15° give account of typical human face variations.

It is important to remark that the acquisition device does not require any special installation, and therefore it can be used at different places. Thus, the whole acquisition process can be carried out at different locations (street environment, laboratory environment, mall environment, etc.). In our case we use at least two different locations for each person, one indoor (laboratory with windows), and one outdoors (gardens inside our school's campus). See example in fig. 1b.

Background images for each place, camera-height, and yaw-pitch angle combination are taken with the acquisition device, in order to be able to compose the final images to be shown to the agent.

In fig. 2 are shown some examples of images taken with the device.

3.2 Database Description

Face images of 50 persons compose the database. In each case 726 registered face images ($121 \times 3 \times 2$) are stored. The yaw angle range is -120° to 120° , with a resolution of 2° , which gives 121 images. For each different yaw, 3 different pitch angles are considered. For each yaw-pitch combination, indoor and outdoor images are taken. In

addition, background images corresponding to the different yaw-pitch angles, place and camera-height combinations are also stored.




Fig. 1. (a) Diagram of the image acquisition system. (b) The system operating in outdoors.




Fig. 2. Examples of images taken using the device in indoors/outdoors first/second row.

3.3 Virtual Scenario Description and Agent Positioning

The scenario contains real face images of N persons. An observing agent, the one with the ability to recognize faces, has the possibility of navigating and making observations inside the scenario. Considering that the goal of this tool is to test the

recognition abilities of the agent navigation is simplified: the agent is placed at a fixed distance of 100 cm in front of each person by the system. Persons are chosen in a random order. In this first version of the system, the agent's camera and the observed face are at the same height, and the agent cannot move its head independently of the body. The following variations in the agent's relative position and viewpoint are incorporated before the agent starts to recognize person i :

- The pose of the agent is randomly modified in $\Delta x_i, \Delta y_i, \Delta \theta_i$. The maximal variation in each axis, $(\Delta x_{\max}, \Delta y_{\max}, \Delta \theta_{\max})$, is a simulation parameter.
- The face of person i is randomly rotated in θ_i^y (yaw angle), θ_i^p (pitch angle), and θ_i^r (roll angle). The maximal allowed rotation value in each axis, $(\theta_{\max}^y, \theta_{\max}^p, \theta_{\max}^r)$, is a simulation parameter.

After the relative position and orientation between the agent and the observed face are fixed, the simulator generates the corresponding observations (i.e. images) to the agent. This image generation process, more than a rendering process is a image composition process, in which real face and background images acquired using the device described in section 3.2, are used. The out-of-plane rotations are restricted to the available face images in the sagittal and lateral planes, while there are no restrictions for the in-plane rotations. In addition, the system selects at random whether person i is observed under indoor or outdoor illumination conditions.

The agent analyzes the generated images to detect and recognize human faces. Depending on the results of this analysis, the agent can change its relative pose using the following functions:

- $Move(\Delta x, \Delta y)$: It changes the relative position of the agent in x and y . It is considered that the agent has the ability to perform omnidirectional movements.
- $Turn(\Delta \theta)$: The agent turns in $\Delta \theta$. The angle's sign gives the turn direction.

Every time that the agent changes its pose, the simulator generates/composes the corresponding images. For instance, fig. 3 shows a given sequence of agent poses, and the corresponding images composed by the simulator. When the agent decides that it already knows the identity of the person, or that it cannot determine it, it sends this information to the simulation tool. Then, the simulator place the agent in front of the next person, and the whole process is repeated. If there is no next person, then the simulation finishes, and the simulation tool write a log file with the statistics of the recognition process.

3.4 Testing Methodology

In order to recognize faces properly, the agent needs to have the following functionalities: (i) *Face Detection*. The agent detects a face (i.e. the face region) in a given image; (ii) *Face Pose Estimation*. The agent estimates the face's angular pose in the lateral, sagittal and coronal plane; (iii) *Active Vision*. Using information about the detected face and its pose, and other information observed in the input images, the agent can take actions in order to change the viewpoint of the sensor for improving

face's perception; and (iv) *Face Recognition*. The identity of the person contained in the face image is determined.

In the testing tool these functionalities are implemented by the *DetectFace*, *EstimateFaceAngularPose*, *ImproveAgentPose*, and *RecognizeFace* functions (see fig. 4). The face recognition system under analysis should have at least the *RecognizeFace* function; having the other functions is optional. The testing tool can provide the functions that the face recognition system is not including. In case the testing tool is providing the *DetectFace* and *EstimateFaceAngularPose*, the face detection and face pose estimation accuracy can be fully controlled (the simulator knows the ground truth). They are simulation parameters to be defined in the testing protocol.

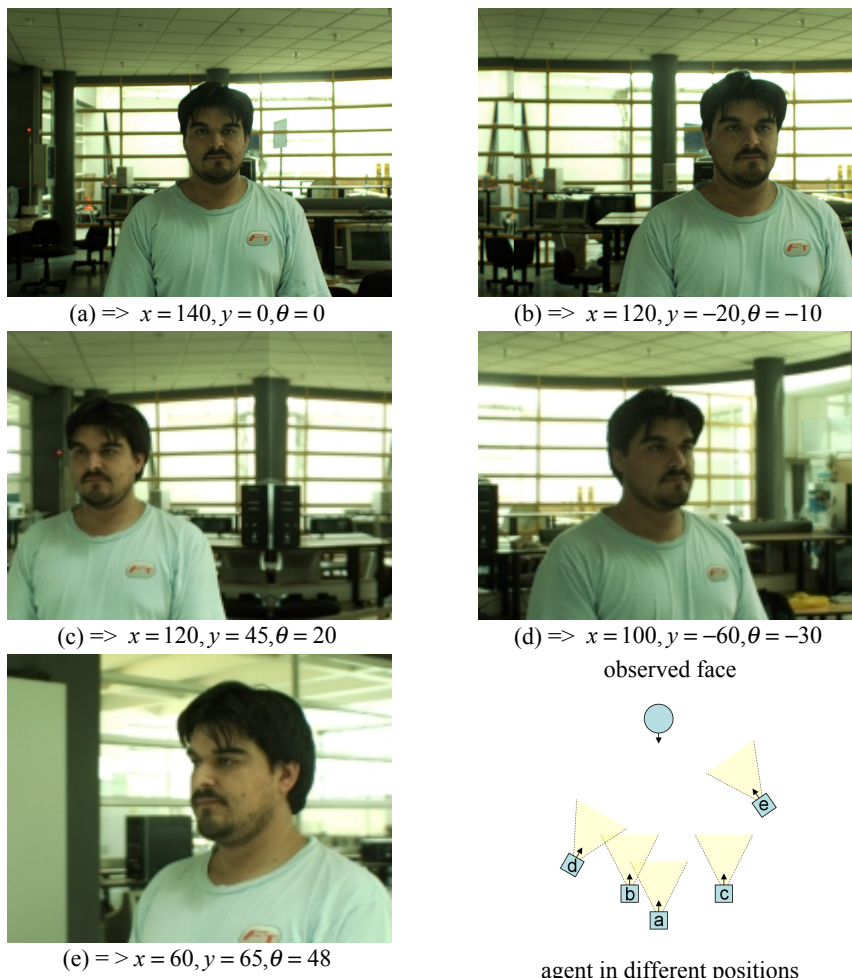


Fig. 3. Example of agent's positioning and the images/observations composed by the simulator. The agent is located in (a), and it moves to positions (b)-(e). In each case the input images are shown.

The simulation tool allows using the following modes:

- *Mode 1 - Recognition using Gallery*: The simulation tool generates a face gallery before the recognition process starts. The gallery contains one image of each person to be recognized. The gallery's images are frontal pictures (no rotations in any plane), taken under indoor illumination conditions. This is the standard operation mode, whose pseudo algorithm is shown in figure 4.

- *Mode 2 - Recognition without using Gallery*: There is no gallery. The agent needs to cross two times the virtual scenario. In the first round, it should create the database online. In the second round, the gallery is used to recognize the persons. In both rounds, the agent see the person's faces at variable distance and angles, in indoor or outdoor illumination conditions. The persons pose and the illumination conditions are randomly chosen.

In each of the two described modes it can be activated the option $-m$, which allows observing multiple persons in some images. In this case, the persons were previously scanned together by the image acquisition system.

4 Results

In order to obtain a first validation of the applicability of the testing tool, three unsupervised face recognition methods are compared. In the reported experiments, face detection, face pose estimation, and active vision are provided by the testing tool.

4.1 Face Recognition Methods

Three local-matching face recognition methods are implemented: histograms of LBP (Local Binary Patterns) features [22], Gabor-Jet features with Borda count classifiers [23], and histograms of WLD (Weber Local Descriptor) features. The first two methods have shown a very good performance in comparative studies of face recognition systems [10][23]. The third method is being proposed here, and it is based in the recently proposed WLD feature [24]. In all cases, the methods' parameters are adapted/adjusted using standard face datasets, and not using the face images that the testing tool includes.

Following the results reported in [10], two different flavors of the histograms of LBP features method are implemented, one using the histogram intersection (HI) similarity measure, and one using the Chi square (XS) measure. In both cases face images are scaled to 81x150 pixels and divided into 40 regions to compute the LBP histograms. The two implemented face recognitions systems are called LBP-HI-40 and LBP-XS-40. The implemented Gabor-based method uses 5 scales and 8 orientations, and face images scale to 122x225 pixels, as reported in [10].

Finally, in the case of the WLD based method, after extensive experimentation using the FERET, BioID and LFW databases, the following parameters were selected: histogram intersection and Chi square similarity measures, face images scaled to 93x173 pixels and divided into 40 regions to compute the WLD histograms, 2 dominant orientations ($T=2$), and 26 cells in each orientation ($C=26$).


```

Initialization:
    SetMaxVariationAgentInitialPosition ( $\Delta x_{\max}$ ,  $\Delta y_{\max}$ ,  $\Delta \theta_{\max}$ );
    SetMaxVariationFaceRotationAngles ( $\theta_{\max}^y$ ,  $\theta_{\max}^p$ ,  $\theta_{\max}^r$ );
    num_recognized_faces=num_false_positives=0;
Testing:
    for (i=0;i<N;i++)
        SetRobotInitialPose();
        SetFaceInitialPose();
        SetIndoorOutdoorIllumination();
        currentImage =GetImage();
        id=RecognizePerson(currentImage);
        if (id==GetPersonID(i))
            num_recognized_faces+=1;
        else if (id!= NO_IDENTIFICATION)
            num_false_positives+=1;
    StoreStatistics(num_recognized_faces, num_false_positives);
end;

Recognition:
    RecognizePerson(image)
        while(1)
            if((face=DetectFace(image))==NO_IDENTIFICATION)
                return(NO_IDENTIFICATION);
            faceAngularPose=EstimateFaceAngularPose(image);
            if(face.size<MIN_SIZE OR |faceAngularPose.yaw|>MIN_YAW OR
                |faceAngularPose.pitch|>MIN_PITCH)
                ImproveAgentPose(face.position, face.size,faceAngularPose);
                image=GetImage();
            else
                result=RecognizeFace(face)
                if(result.confidence<threshold)
                    return(NO_IDENTIFICATION)
                else
                    return(result.id)

```

Fig. 4. Pseudo algorithm of testing procedure in mode 1 (recognition using gallery).

4.2 Recognition Results

In a first set of experiments, the recognition rate of the different methods is compared under different viewpoint conditions; the yaw angle of the observed faces is uniformly selected (random value) in the range $-/+ \theta_{\max}^y$. The other simulation parameters are kept unchanged ($\Delta x_{\max} = \Delta y_{\max} = \Delta \theta_{\max} = \theta_{\max}^p = \theta_{\max}^r = 0$). In the experiments no active vision mechanisms are used, and a face detection rate of 100% is considered. Table 1 shows the obtained results. Main conclusions of these experiments are: (i) LBP based methods that use the Chi square similarity measure are

more robust to yaw rotations than Gabor and WLD based methods, and (ii) all methods are robust to yaw rotation in the range $\pm 30^\circ$.

Table 1. Top-1 recognition rates under different maximal yaw angles of the observed face (θ_{\max}^y). The other parameters are not varied ($\Delta x_{\max} = \Delta y_{\max} = \Delta \theta_{\max} = \theta_{\max}^p = \theta_{\max}^r = 0$).

Method	θ_{\max}^y								
	5°	10°	15°	20°	25°	30°	35°	40°	60°
LBP-HI-40	1.00	1.00	1.00	1.00	0.95	0.95	0.80	0.85	0.55
LBP-XS-40	1.00	1.00	1.00	0.95	0.95	0.95	0.85	0.75	0.30
GJD-BC	1.00	1.00	1.00	0.95	0.85	0.85	0.75	0.80	0.35
WLD-HI-40	1.00	1.00	0.95	0.95	0.90	0.90	0.85	0.70	0.45
WLD-XS-40	1.00	1.00	0.90	0.90	0.95	0.90	0.75	0.70	0.45

Table 2. Top-1 recognition rates under different maximal yaw and pitch angles of the observed face ($\theta_{\max}^y, \theta_{\max}^p$), different maximal agent's positioning errors ($\Delta x_{\max}, \Delta y_{\max}$), and variable face

pose estimation error (pe).

Method	$\theta_{\max}^y = \pm 45$	$\theta_{\max}^y = \pm 45$	$\theta_{\max}^y = \pm 45$	$\theta_{\max}^y = \pm 45$	$\theta_{\max}^y = \pm 45$	$\theta_{\max}^y = \pm 45$
	$\theta_{\max}^p = 0$	$\theta_{\max}^p = 0$	$\theta_{\max}^p = \pm 15$	$\theta_{\max}^p = \pm 15$	$\theta_{\max}^p = \pm 15$	$\theta_{\max}^p = \pm 15$
	$\Delta x_{\max} = 20$	$\Delta x_{\max} = 40$	$\Delta x_{\max} = 20$	$\Delta x_{\max} = 40$	$\Delta x_{\max} = 20$	$\Delta x_{\max} = 40$
	$\Delta y_{\max} = 20$	$\Delta y_{\max} = 40$	$\Delta y_{\max} = 20$	$\Delta y_{\max} = 40$	$\Delta y_{\max} = 20$	$\Delta y_{\max} = 40$
	$pe = 40\%$	$pe = 40\%$	$pe = 40\%$	$pe = 40\%$	$pe = 80\%$	$pe = 80\%$
LBP-HI-40	0.85	0.85	0.80	0.75	0.75	0.70
LBP-XS-40	0.85	0.85	0.80	0.80	0.80	0.75
GJD-BC	0.85	0.80	0.75	0.70	0.70	0.65
WLD-HI-40	0.80	0.85	0.70	0.65	0.70	0.65
WLD-XS-40	0.80	0.85	0.70	0.65	0.70	0.65

In a second set of experiments, the recognition rate of the different methods is compared under more uncontrolled conditions:

- The yaw angle of the observed faces is uniformly selected (random value) in the range $\pm 45^\circ$, and the pitch angle in the range $\pm 15^\circ$. The roll angle is not modified ($\theta_{\max}^r = 0$).
- The position of the observer agent is modified in each axis, by a random value uniformly selected in the range ± 20 or ± 40 centimeters. The agent is not rotated ($\Delta \theta_{\max} = 0$).

The following face detection and pose estimation conditions are considered: (i) Face detection rate of 80% with no false positives, (ii) Face pose estimation with an error, pe , uniformly selected (random value) in the range $\pm 40\%$ or $\pm 80\%$ of the estimated value, and (iii) Active vision mechanisms as shown in the procedure of fig. 4.

Table 2 shows the obtained results. Main conclusions of these experiments are: (i) LBP based methods are more robust to the defined uncontrolled conditions than Gabor and WLD based methods, (ii) the agent's initial position error has a low influence on the final performance of the recognition systems, (iii) a maximal error of

+/-15° in the pitch angle reduces in ~5% the face recognition rate, and (iv) a pose estimation error increase from 40% to 80% reduces in ~5% the recognition rate.

5 Conclusions and Projections

In this article, a tool for testing face recognition systems under uncontrolled conditions is proposed. The testing tool combines the use of a simulator with real face and background images taken under real-world conditions. Inside the simulated environment, an observing agent can navigate and observe the real face images, at different distances, angles and with indoor or outdoor illumination. During the face recognition process, the agent can actively change its viewpoint and relative distance to the faces in order to improve the recognition results. The simulation tool provides all navigation and positioning functionalities to the agent, except the ones related with the detection, alignment and recognition of faces.

The applicability of the proposed tool is validated in the comparison of three state of the art face recognition methods, histograms of LBP features, Gabor-Jet features with Borda count classifiers, and histograms of WLD features.

In order to be able to share the use of the proposed tool with other researchers of the face recognition community, the following procedure will be implemented:

1. A DLL of the testing tool with a sample of the database containing the images of 10 individuals will be distributed upon request³. The DLL will include a Visual Studio project with an example of use. In a second stage, a Linux library will be also provided. The goal of this DLL is that researchers can make parameters' adjustment and preliminary tests of their face recognition methods.
2. In order to carry out tests using the complete database, researchers will submit a compiled version of their face recognition method, linked to the provided DLL. After testing, results will be sent back automatically.

We are currently writing a technical report where the outlined procedure will be described in detail, as well as the conditions of use of the tool. We are also implementing a website to manage the described procedure.

References

1. W. Zhao, R. Chellappa, A. Rosenfeld, P.J. Phillips, Face Recognition: A Literature Survey, *ACM Computing Surveys*, 2003, pp. 399-458.
2. X. Tan, S. Chen, Z.-H. Zhou, and F. Zhang, Face recognition from a single image per person: A survey, *Pattern Recognition*, Vol. 39, pp. 1725-1745, 2006.
3. R. Chellappa, C.L. Wilson, S. Sirohey, Human and Machine Recognition of Faces: A Survey, *Proceedings of the IEEE*, Vol. 83, Issue 5, May 1995, pp. 705-740.
4. A. F. Abate, M. Nappi, D. Riccio, G. Sabatino, 2D and 3D face recognition: A survey, *Pattern Recognition Letters*, Vol. 28, pp. 1885-1906, 2007.
5. Face Recognition Home Page: <http://www.face-rec.org/databases/>

³ The full database cannot be made available because of its size.

6. Call for Papers Special Issue in Real-World Face Recognition, IEEE Trans. on PAMI : http://www.eecs.northwestern.edu/~ganghua/ghweb/CFP_TPAMI_FR.htm
7. *Faces in Real-Life Images Workshop*, Oct. 17th 2009, ECCV 2008: <http://hal.inria.fr/REALFACES2008/en>
8. *Labeled Faces in the Wild Database*: <http://vis-www.cs.umass.edu/lfw/index.html>
9. G.B. Huang, M. Ramesh, T. Berg, and E. Learned-Miller, Labeled Faces in the Wild: A Database for Studying Face Recognition in Unconstrained Environments. University of Massachusetts, Amherst, Technical Report 07-49, Oct. 2007.
10. J. Ruiz-del-Solar, R. Verschae, M. Correa (2009). Recognition of Faces in Unconstrained Environments: A Comparative Study. *EURASIP Journal on Advances in Signal Processing (Recent Advances in Biometric Systems: A Signal Processing Perspective)*, Vol. 2009, Article ID 184617, 19 pages.
11. M. Jones (2009). Face Recognition: Where We Are and Where To Go From Here, Mitsubishi Electric Research Laboratories Technical Report TR2009-023, June 2009.
12. Face Recognition Grand Challenge, Official website. Available in <http://www.frvt.org/FRGC/>
13. P. J. Phillips, H. Wechsler, J. Huang and P. Rauss, The FERET database and evaluation procedure for face recognition algorithms, *Image and Vision Computing J.*, Vol. 16, no. 5, pp. 295-306, 1998.
14. P. Phillips, P. Flynn, T. Scruggs, K. Bowyer, J. Chang, K. Hoffman, J. Marques, J. Min, and W. Worek, Overview of the Face Recognition Grand Challenge, Proc. of the IEEE Conf. Computer Vision and Pattern Recognition – CVPR 2005, Vol. 1, p. 947-954.
15. R. Gross, Face Databases, in *Handbook of Face Recognition*, S. Li and A.K Jain (Eds.), Springer-Verlag, pp. 301-327, 2005.
16. Yale University Face Image Database, publicly available in <http://cvc.yale.edu/projects/yalefaces/yalefaces.html>
17. BioID Face Database, publicly available in <http://www.humanscan.de/support/downloads/facedb.php>
18. AR Face Database public site http://cobweb.ecn.purdue.edu/~aleix/aleix_face_DB.html
19. P.J. Flynn, K.W. Bowyer, and P.J. Phillips (2003). Assessment of time dependency in face recognition: An initial study, *Audio and Video-Based Biometric Person Authentication*, 44-51.
20. Yale Face Database B. Available in <http://cvc.yale.edu/projects/yalefacesB/yalefacesB.html>
21. PIE Database. Basic information in: http://www.ri.cmu.edu/projects/project_418.html
22. T. Ahonen, A. Hadid, and M. Pietikainen, Face Description with Local Binary Patterns: Application to Face Recognition, *IEEE Trans. on Patt. Analysis and Machine Intell.*, Vol. 28, No. 12, pp. 2037-2041, Dec. 2006.
23. J. Zou, Q. Ji, G. Nagy, A Comparative Study of Local Matching Approach for Face Recognition, *IEEE Trans. on Image Processing*, Vol. 16, Issue 10, Oct. 2007, pp. 2617-2628.
24. J. Chen, S. Shan, C. He, G. Zhao, M. Pietikäinen, X. Chen, and W. Gao, WLD: A Robust Local Image Descriptor, *IEEE Trans. on Patt. Analysis and Machine Intell.*, TPAMI-2008-09-0620 (in press).