TABULA RASA

Trusted Biometrics under Spoofing Attacks

http://www.tabularasa-euproject.org/

Funded under the 7th FP (Seventh Framework Programme)
Theme ICT-2009.1.4
[Trustworthy Information and Communication Technologies]

D2.4: ICAO biometric databases of spoofing attacks

Due date: 29/02/2012
Submission date: 29/02/2012
Project start date: 01/11/2010
Duration: 42 months
WP Manager: Javier Acedo (STARLAB)
Revision: 1

Author(s): Andre Anjos (IDIAP), Julian Fierrez (UAM), Stan Li and Dong Yi (CASIA), Neslihan Kose (EURECOM), Nguyen Eric (MORPHO)

Project funded by the European Commission in the 7th Framework Programme (2008-2010)		
Dissemination Level		
PU	Public	Yes
RE	Restricted to a group specified by the consortium (includes Commission Services)	No
CO	Confidential, only for members of the consortium (includes Commission Services)	No

D2.4: ICAO biometric databases of spoofing attacks

Abstract:

This document describes the different biometric databases of spoofing attacks collected for the ICAO modalities. These modalities include fingerprint, face (2D, 3D, multi-spectral) and iris that coherent with the description in D2.3, and will be used for algorithms evaluation (WP3) and countermeasures development (WP4). For each modality, a brief spoofing attack associated with the database will be described firstly. Then, a detail description of the data and the protocol of how to use the database will be given. For more information on the spoofing attacks, please refer to D2.3 Specifications Of Spoofing Attacks.

Contents

1 Introduction

The main objective of TABULA RASA is to find countermeasures of different biometric systems to possible spoofing attacks. In order to do so, the first thing to do is to define the possible spoofing attacks, or in other words, to find the potential weaknesses of existing biometric systems. Once the spoofing attacks have been defined in "D2.3: Specification of Spoofing attacks", a database containing those spoofing attacks was planned to be recorded. Following the ICAO biometric standards, we consider separately different biometrics. Fingerprint, Face, and Iris are referred as ICAO biometrics, and others are referred as non-ICAO biometrics. This deliverable describes those different databases for the ICAO biometric modalities, which are Fingerprint, 2D-Face, 3D-Face, Multi-spectral Face, and Iris.

The structure of each section is as follows. For each modality we have defined 4 subsections: spoofing attacks description (which is a summary of the attacks defined in D2.3), database description, the protocol for licit biometric transactions and the protocol for spoofing attacks. Using the databases and the corresponding protocols, we can evaluate the baseline system under spoofing attacks, and the performance of countermeasures.

2 2D Face Databases of Spoofing Attacks

2.1 Spoofing Attacks Description

The 2D face spoofing attack database allows users to measure the effectiveness of spoofing attacks or counter-measures to 2D face recognition systems. It is composed of two sets of data: real-access and attacks. Real-accesses are used to establish reference performance figures for recognition systems. Attacks can be used to model spoofing classifiers or measure the impact of spoofing to existing baselines.

In the reminder of this section we describe technical details of how such database was created and its contents. We close such a description by presenting protocols for database usage in different scenarios that occur during anti-spoofing algorithm development: definition of baseline recognition performances, sensitiveness to spoofing attacks and the development of counter-measures.

2.2 Database Description

The 2D face spoofing attack database consists of $1,300$ video clips comprising of real-accesses or photo and video attack attempts to different 50 identities, under different lighting conditions.

The data is split into 3 sub-groups comprising:

1. Training data ("train"), to be used for training your anti-spoof classifier;
2. Development data ("devel"), to be used for threshold estimation;
3. Test data ("test"), with which to report error figures;

Clients that appear in one of the data sets (train, devel or test) do not appear in any other set. Figure 1 provides some example snapshots of attack videos.

2.3 Setup

All videos are generated by either having a (real) client trying to access a laptop through a built-in webcam (see Figure 2) or by displaying a photo or a video recording of the same client for at least 9 seconds. The webcam produces color videos with a resolution of 320 pixels (width) by 240 pixels (height). The movies were recorded on a Macbook laptop using the QuickTime framework (codec: Motion JPEG) and saved into ".mov" files. The frame rate is about 25 Hz. Besides the native support on Apple computers, these files are *easily* readable using mplayer, ffmpeg or any other video utilities available under Linux or MS Windows systems.

Real client accesses as well as data collected for the attacks are taken under two different lighting conditions:

Figure 1: Example attacks in different scenarios and with different lighting conditions. On the top row, attacks in the *controlled* scenario. On the bottom, attacks with samples from the *adverse* scenario. Columns from left to right show examples of hard-print, mobile phone and high-resolution screen attacks.

Figure 2: Setup used for the acquisition of real-accesses for the 2D face spoofing database.

- **controlled**: The office light was turned on, blinds are down, background is homogeneous;
- **adverse**: Blinds up, more complex background, office lights are out.

Real-accesses are recorded directly with the acquisition system webcam. Each user captures 6 videos:

- 2 enrollment videos, for creating user models
- 2 authentication videos, for testing user access under controlled conditions;
- 2 authentication videos, for testing user access under adverse conditions.

To produce the attacks, high-resolution photos and videos from each client were taken under the same conditions as in their authentication sessions, using a Canon PowerShot SX150 IS camera, which records both 12.1M pixel photographs and 720p high-definition video clips. The way to perform the attacks can be divided into two subsets: the first subset is composed of videos generated using a stand to hold the client biometry ("fixed"). For the second set, the attacker holds the device used for the attack with their own hands. In total, 20 attack videos were registered for each client, 10 for each of the attacking modes just described:

- 4× mobile attacks using an iPhone 3GS screen (with resolution 480x320 pixels) displaying:
  - 1× mobile photo/controlled
  - 1× mobile photo/adverse
  - 1× mobile video/controlled
  - 1× mobile video/adverse
- 4× high-resolution screen attacks using an iPad (first generation, with a screen resolution of 1024x768 pixels) displaying:
  - 1× high-resolution photo/controlled
  - 1× high-resolution photo/adverse
  - 1× high-resolution video/controlled
  - 1× high-resolution video/adverse
- 2× hard-copy print attacks (produced on a Triumph-Adler DCC 2520 color laser printer) occupying the whole available printing surface on A4 paper for the following samples:
  - 1× high-resolution print of photo/controlled

– 1× high-resolution print of photo/adverse

The real-access (authentication) and attack videos were then divided in the following way:

- **Training set**: contains 60 real-accesses and 300 attacks under different lighting conditions;
- **Development set**: contains 60 real-accesses and 300 attacks under different lighting conditions;
- **Test set**: contains 80 real-accesses and 400 attacks under different lighting conditions.

The 100 enrollment videos (2 per client) are put aside for exclusively studying the baseline performance of face recognition systems as explained next.

2.3.1 Face Locations

We also provide face locations automatically annotated by a cascade of classifiers based on a variant of Local Binary Patterns (LBP) referred as Modified Census Transform (MCT) [2]. The automatic face localisation procedure works in more than 99% of the total number of frames acquired. This means that less than 1% of the total set of frames for all videos do not possess annotated faces. User algorithms must account for this fact.

2.4 Protocol for Licit Biometric Transactions

It is possible to measure the performance of baseline face recognition systems using specific data available at the 2D Face spoofing database. This step reports how well such a system behaves in normal operation mode, with real-access attempts. In this setting the system is tested for how well it can recognize real users authenticating against their templates and how well it can reject real users authenticating against other users templates. The performance is measured objectively by observing the rate of users rejected when authenticating against their own template (False Rejection Rate - FRR) and by the rate of users accepted when authenticating against someone else's template (False Acceptance Rate - FAR). In other words we establish the baseline recognition performance of the face recognition system using exclusively real-access (enrollment attempts and authentication tries).

We describe steps to be followed for building and tunning the TABULA RASA baseline 2D face recognition system and to assess its performance with the data available in this spoofing database. Other systems may be adapted from these instructions for as long as testing and tunning are performed with the disjoint sets indicated by this protocol.

2.4.1 Baseline Recognition Performance

The TABULA RASA 2D face baseline algorithm is based on generative models and requires the construction of an universal background model (UBM) and the tunning of a global threshold for the separation of scores generated when matching new user imagery against its pre-computed template. The UBM should be exclusively built from the training-set client enrollment videos. If a global separation threshold is required, it should be defined exclusively using the development-set clients. Details are as follows:

1. Universal Background Model (UBM): To generate the UBM, select the training-set client videos from the enrollment videos. There should be 2 per client, which means one gets 30 videos, each with 375 frames to create such a model;

2. Client models: To generate client models, use the enrollment data for clients at the development and test groups. There should be 2 videos per client (one for each light condition) once more. At the end of the enrollment procedure, the development set must have 1 model for each of the 15 clients available in that set. Similarly, for the test set, 1 model for each of the 20 clients available.

To determine the baseline face recognition performance, generate scores **exhaustively** for all videos from the development and test **real-accesses** respectively, testing access from every identity in the set against all other models in the set, but **without** intermixing across development and test sets. The scores generated against matched client videos and models (within the subset, i.e. development or test) should be considered true client accesses, while all others, impostors. By varying the classification threshold on the test set, one produces the recognition performance of the system. In this context the FAR (False Acceptance Rate) should be considered the rate of impostors that are wrongly classified by the system as true-claimants. The FRR (False Rejection Rate) is the rate of true claimants that get falsely classified as impostors.

If a single number is required, a threshold should be set using the development data and then applied to the test scores and a number reported from the latter.

2.5 Protocols for Spoofing Attacks

When considering spoofing attacks, the enrollment is achieved using real-accesses, as defined in the previous section. Tests are carried out using both real-accesses and attacks to corresponding identities only. A successful attack is accomplished when the system confuses a spoofing attempt with the corresponding matched user template. In this mode, the FAR corresponds to the rate of attacks that are accepted by the system when spoofed. The FRR corresponds to the rate of real-access attempts that are incorrectly dismissed by the system as attacks.

The database can be split into 6 different protocols according to the type of device used to generate the attack: print, mobile (phone), high-definition (tablet), photo, video

Table 1: Distribution of videos per attack-protocol in the 2D Face spoofing database.

	Hand-Attacks			**Fixed-Support**			**All Supports**		
Protocol	train	dev	test	train	dev	test	train	dev	test
Print	30	30	40	30	30	40	60	60	80
Mobile	60	60	80	60	60	80	120	120	160
Highdef	60	60	80	60	60	80	120	120	160
Photo	90	90	120	90	90	120	180	180	240
Video	60	60	80	60	60	80	120	120	160
Grandtest	150	150	200	150	150	200	300	300	400

or grand test (all types). All attack protocols are composed of the same number of real-access authentication and enrollment data as presented before. Further subsetting can be achieved on the top of the previous 6 protocols by classifying attacks as performed by the attacker bare hands or using a fixed support. This classification scheme makes-up a total of 18 protocols that can be used for studying the performance of counter-measures to 2D face spoofing attacks. Table 3 details the amount of video clips in each protocol.

Table 2: The composition the 3D Face spoofing database.

Type of acquisition	Number of subjects	Number of scans per subject
Black&White masks (real 3D printing)	20	10
Real faces (3D acquisitions)	20	10

3 3D Face Databases of Spoofing Attacks

3.1 Spoofing Attacks Description

In order to measure the impact of spoofing attacks and efficiency of counter-measures on the performance of 3D face recognition systems, a 3D face spoofing attacks database is created. The 3D face baseline system is using face shape characteristics to recognize a user. The texture of the face or lighting is not taken into account in the system. The acquisition is made with a 3D scanner which will reconstruct the face shape of the user. Therefore, the spoofing attacks are acquired from people who are wearing 3D masks.

3.2 Database Description

The 3D face database of spoofing attacks consists of 200 acquisitions of 16 black and white masks coming from 3D scanner. The database includes 20 people, among which 16 people's faces have been used for mask manufacturing. The remaining 4 people do not have masks but they are also involved in the mask acquisitions by wearing the masks of other people. Each subject involved in the mask acquisition has 10 scans. As 3D scanners are not common and the acquisition process is really sensitive to movement, it can be done only with the cooperation of the user. To simulate the practical situations, all acquisitions are done under unconstrained environments. In the process of acquisition, the angle of face is variable and the lighting is not necessary uniform.

The table 2 summarizes the database.

Note that in mask acquisitions, number of subjects involved is 20 whereas the number of masks used for the mask acquisition is 16. Furthermore even the subject wear his/her own mask to attack the system, it will be considered as a mask attack.

The database is split into 3 sub-groups comprising:

1. Training set: contains 70 real-accesses and 70 attacks;
2. Development set: contains 60 real-accesses and 60 attacks;
3. Test set: contains 70 real-accesses and 70 attacks.

Clients that appear in one of the data sets (train, devel or test) do not appear in any other set.

Figure 3: Left: a 3D model generated from point cloud, Right: its corresponding back&white mask made by a 3D printer.

3.3 Spoofing attacks manufacturing

3.3.1 Black and white mask coming from 3D scanner

First, a 3D acquisition is performed with a 3D scanner. The 3D acquisition is a point cloud, which cannot be printed as there is no surface, therefore it is mandatory to generate a mesh from the point cloud. This mesh is extrapolation of the acquisition into a polygon 3D model. Then, point cloud becomes a 3D structure of polygons:

Then, this polygons 3D model is sent to a 3D printer. Figure 3 shows an example of the result.

3.3.2 Colored Masks Coming from Extrapolation of 2D Pictures

With this method, 2 pictures of the face of the end-user are needed: picture of the front face and picture of one side. Then a 3D model is calculated from the 2D pictures, knowing that it is based on the ethnic shape. Finally it is sent to 3D printer and a mask is manufactured. The 3D face system selected for this project is tested with some colored mask samples however the system is robust to colored mask attacks. Therefore, colored masks are not involved in the database.

3.3.3 2D Pictures Printed on T-shirt

With this method, we take a high definition photo of an end-user and printed it on a white T-shirt. Then the t-shirt is put on a face of someone who is different of the printed face and a 3D acquisition is done. Again the 3D face system selected for this project is tested with some t-shirt attack samples however the system is robust to t-shirt attacks. Therefore, t-shirt attacks are not involved in the database.

3.4 Protocol for Licit Biometric Transactions

This database includes acquisition of real people and acquisition of masks. It will be possible to measure the biometric performance of the algorithm with a baseline, then to measure the performance with the spoofing attacks. Using the real access data in the 3D face spoofing database, we can measure the performance of baseline face recognition systems. This step reports how well such a system behaves in normal operation mode. The steps for evaluation a 3D face recognition system are referred to D3.1. In this part, each acquisition will be compared with all other acquisition in the database.

In licit biometric transactions part, the system is tested for how well it can recognize real users authenticating against their templates and how well it can reject real users authenticating against other users templates. This means that the FAR (False Acceptance Rate) should be considered as the rate of impostors that are wrongly classified by the system as true-claimants and the FRR (False Rejection Rate) is the rate of true claimants that get wrongly classified as impostors.

The real face scans in the database is used to evaluate the performance of the baseline face recognition system. In this part, we describe how to use the database to do training, enrollment, and recognition.

For the experiment, the database is split into 3 groups:

1. Training set: real face accesses of the 7 subjects, 7×10=70;
2. Development set: real face accesses of the other 6 subjects, 6×10=60;
3. Test set: real face accesses of the other 7 subjects, 7×10=70.

The 3D Face recognition system will use the real access images in the test set to evaluate the baseline performance. Since in the mask acquisitions, the number of subjects involved (20 subjects) is not equal to the number of masks used (16 masks), the test set subjects should be selected from the subjects whose masks are manufactured for this database. Access from every identity in the test set will be tested against all other models in the test set. The scores generated against matched clients should be considered true client accesses, while all others are impostors.

The Detection-Error Trade-of (DET) curve is used to illustrate the baseline performance of the 3D face recognition system.

3.5 Protocol for Spoofing Attacks

In case of spoofing attacks, the enrollment is again achieved using real-accesses, as similar to the work explained in the previous section and this time tests are carried out using both real-access attempts and also attacks to corresponding identities only. A successful attack is accomplished when the system confuses a spoofing attempt with the corresponding matched user template.

The table 3 summarizes the number of mask acquisitions which will be used in training, development and test sets. One scenario can be evaluated to illustrate the influence of

Table 3: The composition the 3D Face spoofing database.

Protocol	Train	Development	Test
Black&White masks	7×10	6×10	7×10

spoofing attacks to recognition performance. Real-accesses of the 7 subjects in the test set are used in the gallery and tests are carried out using both real accesses and black&white mask attacks to corresponding identities.

In this part, the FAR corresponds to the rate of attacks that are accepted by the system when spoofed. The FRR corresponds to the rate of real-access attempts that are incorrectly dismissed by the system as attacks. The Detection-Error Trade-of (DET) curves are to be utilized, which describes the relationship between false detection rate and false rejection rate.

4 Multi-spectral Face Databases of Spoofing Attacks

4.1 Spoofing Attacks Description

In a multi-spectral face recognition system, active illuminations of certain spectrums are provided for imaging, which effectively erases the environmental effects.

To effectively attack a multi-spectral face recognition system, the attacker firstly has to capture target subject's face images under the illumination at the same spectrum as the system adopts. Then proper exhibition medium has to be chosen because improper medium will cause extremely strong reflectance. This will innately make fake face attacks difficult, hence helping to strengthening the system's anti-spoofing ability.

For multi-spectral face biometrics, both visible fake faces and multi-spectral fake faces may be potential threatens and therefore should be considered together. Traditionally printed photos which are captured and displayed in visible illumination are just like the one in the top-left of Figure 5. For multi-spectral oriented photos, they have to be captured under the same spectrum as the face recognition systems adopt. Video playbacks are not considered due to the serious specular reflectance. One noticeable point is the photo material used for attacking, because smooth and high quality papers tend to reflect active illumination strongly. Therefore some coarse and less reflective papers have to be used instead, such as newsprint paper, xuan paper, etc.

4.2 Database Description

All face images in the database are acquired by using a self-developed device, which includes a VIS (visual light) camera, a NIR (near infrared, 850nm) camera with some NIR LEDs. VIS and NIR images are synchronized by the system. The imaging device works at a rate of 30 frames per second with the USB 2.0 protocol for 640×480 images. Fig. 4 shows the scenario of database collection.

The database consists the real face images of 100 subjects and their corresponding fake face samples. All subjects are imaged under visible and Near-Infrared (5 images per subject per spectrum) illuminations. For the fake faces, photos are printed using both the visible and NIR face image, named VIS photo and NIR photo respectively. And the printed photos are acquired by the same system described above. We use a kind of coarse paper as printing material, because its relatively rough surface, making the reflectance less strong and the fake face more vivid. Some examples can be seen in Fig. 5.

In summary, the information of the database is: We denote the these kinds of face images as:

- G_{VIS}: Genuine subjects captured by VIS imaging device;
- G_{NIR}: Genuine subjects captured by NIR imaging device;
- VP_{VIS}: VIS photo captured by VIS imaging device;
- VP_{NIR}: VIS photo captured by NIR imaging device;

Figure 4: Setup used for the acquisition of real-accesses for the multi-spectral face spoofing database.

Figure 5: Illustration of multi-spectral face images. From left to right, the three columns correspond to genuine face, VIS photo and NIR photo. Images in the first row are captured under visible illumination, and the second row are captured under 850nm NIR illumination.

Table 4: The information of the multi-spectral face spoofing database.

	Genuine subjects	VIS photo attack	NIR photo attack
Captured by VIS	100 subjects × 5	100 subjects × 5	100 subjects × 5
Captured by NIR	100 subjects × 5	100 subjects × 5	100 subjects × 5

- NP_{VIS}: NIR photo captured by VIS imaging device;
- NP_{NIR}: NIR photo captured by NIR imaging device.

4.3 Protocol for Licit Biometric Transactions

The genuine face images in the database could be used to evaluate the performance of the baseline face recognition system. Here we describe how to use the database to do training, enrollment, and recognition. To illustrate the performance of individual modality (VIS and NIR) and improvement of the multi- spectral fusion, three sub-experiments were conducted:

- VIS vs. VIS
- NIR vs. NIR
- VIS+NIR vs. VIS+NIR

For the three experiments, the database is split into 4 sub-set including:

- Training set: genuine face images of the the first 30 subjects, 30×5=150 pairs;
- Development set: genuine face images of the the following 30 subjects, 30×5=150 pairs;
- Licit Gallery set: 2 genuine face images of the other 40 subjects, 40×2=80 pairs.
- Licit Probe set: 3 genuine face images of the other 40 subjects, 40×3=120 pairs.

The Detection-Error Trade-off (DET) curve is used to illustrate the performance of the system. Since the VIS and NIR images are always captured in pairwise way, their similarity scores could be fused, e.g, we use the score level fusion in D3.1, and achieve good results.

4.4 Protocol for Spoofing Attacks

As same as the experiments in Licit Biometric Transactions, four scenarios will be evaluated to illustrate the influence of spoofing attacks to recognition performance. The first scenario (S1) uses VIS photos to attack the VIS sub-system. The second scenario (S2) uses NIR

photos to attack the NIR sub-system. The third scenario (S3.1) uses the VP_{VIS} and VP_{NIR} photo pair to attack the multi-spectral system, which is the case will be happened in the real world application. The fourth scenario (S3.2) use the NP_{VIS} and NP_{NIR} photo pair to attack the multi-spectral system.

In the attacking scenario, the gallery sets are the same as those in licit transaction, and the probe sets are augmented by the printed photos. In these cases the genuine user enrolls with their faces and the attacker tries to access the system with the corresponding printed photos. A successful attack is accomplished when the system confuses a printed photo with its corresponding genuine face image. The protocol is shown as follows:

- S1 Probe set: Licit Probe set described above and their corresponding printed photo in VP_{VIS};
- S2 Probe set: Licit Probe set described above and their corresponding printed photo in NP_{NIR};
- S3.1 Probe set: Licit Probe set described above and their corresponding printed photo in VP_{VIS} and VP_{NIR};
- S3.2 Probe set: Licit Probe set described above and their corresponding printed photo in NP_{VIS} and NP_{NIR};

The Detection-Error Trade-off (DET) curves are to be utilized, which describes the relationship between false detection rate and false rejection rate. From DET curves, the point where FAR equals FRR is located, and the corresponding value is called the Equal Error Rate (EER), which should also be reported for numerical analysis.

5 Iris Databases of Spoofing Attacks

5.1 Spoofing Attacks Description

For iris modality, two different spoofing attacks are considered:

- **Photo attacks.** A photo-attack consists of displaying a photograph of the attacked identity to the input camera of the iris recognition system. It is easy to obtain face photo of a valid user through internet downloading or simply by capturing them using a concealed/hidden camera. Once a photo is obtained, if it's of sufficient quality, one can print the iris and then present it in front of the iris camera (see Fig. 6). A screen could also be used for presenting the photograph to the camera. Another way to obtain the iris would be to steal the raw iris acquired by an existing iris recognition camera in which the subject being spoofed was already enrolled.

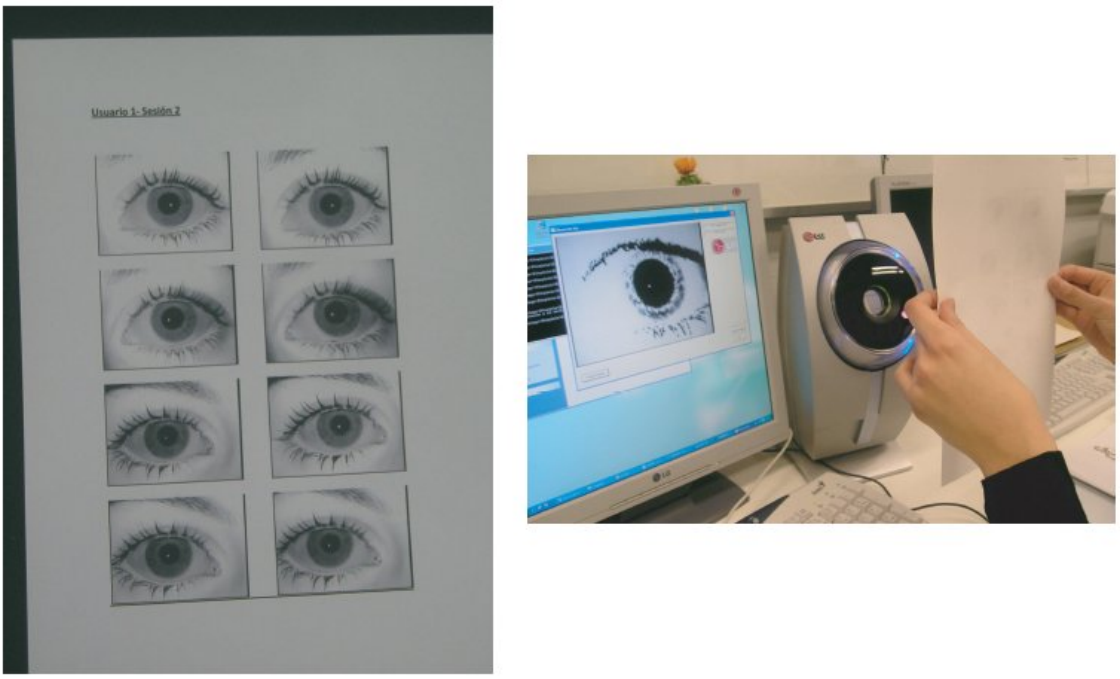

Figure 6: Illustration of images of the acquisition process of a paper fake iris.

- **Attacks with printed contact lenses.** Among the iris counterfeits, a printed contact lens is the most sophisticated. While confronting printed contact lens, the majority anti-spoofing methods will fail to work, e.g., on/off illuminators based methods, frequency analysis, texture analysis, etc. The idea is to print a specific iris pattern (of a subject being attacked) in a contact lens that the attacker will be wearing.

Due to the difficulty in manufacturing printed contact lenses, for this type of attack, only a proof of concept is to be developed in the project. A full database of spoofing attacks has been captured only for the case of photo attacks.

5.2 Database Description

The fake iris database follows the same structure of the original BioSec database [1], from which the fake data is generated. Data consists of 50 subjects $\times$ 2 eyes $\times$ 4 images $\times$ 2 sessions = 800 fake iris images, and its corresponding real images. Acquisition of fake images has been carried out with the same iris camera used in BioSec, a LG IrisAccess EOU3000.

For the experiments, each eye in the database is considered as a different user. In this way, we have two sessions with 4 images each for 100 users (50 subjects $\times$ 2 eyes per subject).

Similarly to what was defined for 2D face databases of spoofing attacks in the present document, here the data is also split into 3 sub-groups comprising:

1. Training data ("train"), to be used for training your anti-spoof classifier. This sub-group corresponds to the first 40 users.

2. Development data ("devel"), to be used for threshold estimation, is constructed with the following 20 users.

3. Test data ("test"), with which to report error figures, corresponds to the last 40 users.

As in the case of 2D face, clients that appear in one of the datasets (train, devel or test) do not appear in any other set. Also, for training and testing both biometric systems and anti-spoofing measures, the same indications for the usage of these subsets that appear for 2D face spoofing databases in the present document should be followed.

5.3 Protocol for Licit Biometric Transactions

In the normal operation mode both the enrollment and the test are carried out with a real iris. This is used as the reference scenario. In this context the FAR (False Acceptance Rate) of the system is defined as the number of times an impostor using his own iris gains access to the system as a genuine user, which can be understood as the robustness of the system against a zero-effort attack. The same way, the FRR (False Rejection Rate) denotes the number of times a genuine user is rejected by the system.

In order to compute the performance of the system in the normal operation mode, the experimental protocol is as follows. For a given user, all the images of the first session are considered as enrolment templates. Genuine matchings are obtained by comparing the templates to the corresponding images of the second session from the same user. Impostor matchings are obtained by comparing the template of the first session to the iris image of the second session.

5.4 Protocol for Spoofing Attacks

In the spoofing scenario the enrollment is performed using a real iris, and tests are carried out with real and fake iriss. In this case the genuine user enrolls with his/her iris and the attacker tries to access the application with the fake iris of the legal user. A successful attack is accomplished when the system confuses a fake iris with its corresponding genuine iris, i.e., the Success Rate of the attack = FAR.

In this spoofing scenario the genuine scores are the same as in the licit transactions scenario. The impostor scores are computed matching all the 4 original samples of each user with its 4 fake samples of the second session.

6 Fingerprint Databases of Spoofing Attacks

6.1 Spoofing Attacks Description

For fingerprint modality, two different spoofing attacks are considered:

- **With cooperation of the user.** In this case the legitimate user takes part in the attack by placing his finger in a small amount of suitable material such as molding silicone; and the impression creates a mold from which artificial fingers can be casted. The artificial fingers are then used in a straightforward way to access the sensor at hand.

- **Without the cooperation of the user.** It is unlikely that in a real-world scenario a user would voluntarily allow to produce an artificial copy of his fingerprints. In this case the gummy fingers can be generated using a similar process to that used in the cooperative scenario, but starting in this place from a latent fingerprint left inadvertently by the user in a surface like a glass. Once the latent fingerprint has been lifted it is printed on to a Printed Circuit Board (PCB) that will serve as mould to produce the artificial finger.

6.2 Database Description

The database is constructed using the index and middle fingers of both hands of 17 users ($17 \times 4 = 68$ different fingers). For each real finger, two fake imitations were generated following each of the procedures explained before (i.e., with and without the user's cooperation). Four samples of each fingerprint (fake and real) were captured in one acquisition session with: *i*) flat optical sensor Biometrika Fx2000 (512 dpi), *ii*) sweeping thermal sensor by Yubee with Atmel's Fingerchip (500 dpi), and *iii*) flat capacitive sensor by Precise Biometrics model Precise 100 SC (500 dpi). Thus, the database comprises 68 fingers $\times 4$ samples $\times 3$ sensors $= 816$ real image samples and as many fake images for each scenario (with and without cooperation). Some good (top rectangle) and bad (bottom rectangle) quality samples (according to the NFIS2 package) of the database are depicted in Fig. 7.

Note that although three different sensors were used for constructing the database of spoofing attacks for fingerprint, the evaluation of attacks and development of countermeasures may not consider the three but only a subset of them.

The database is split into 3 sub-groups comprising:

1. Training data ("train"), to be used for training your anti-spoof classifier. This subgroup corresponds to the first 24 fingers.

2. Development data ("devel"), to be used for threshold estimation, is constructed with the following 20 fingers.

3. Test data ("test"), with which to report error figures, corresponds to the last 24 fingers.

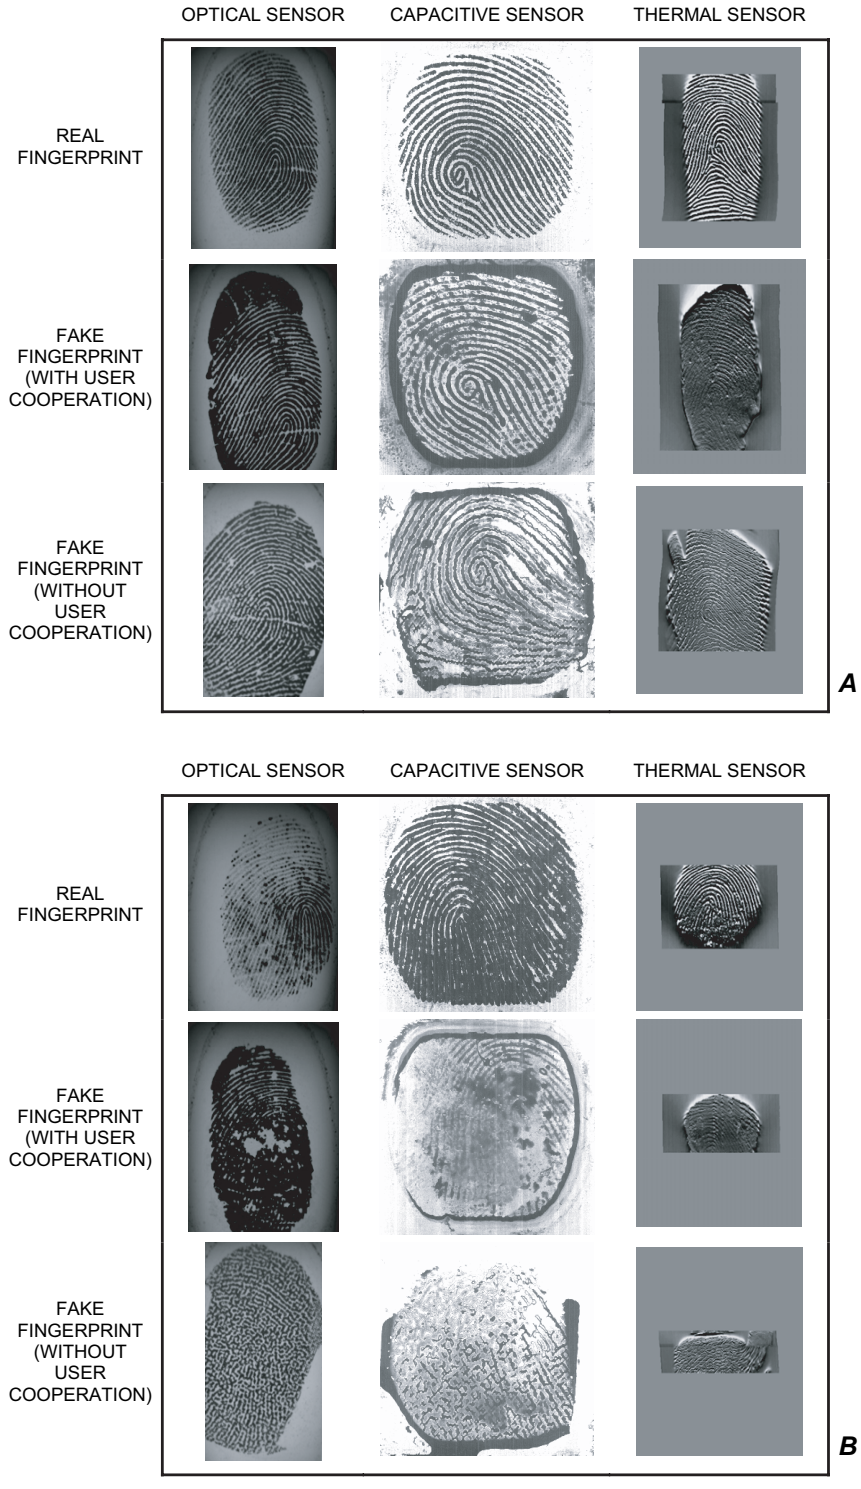

Figure 7: Examples of good (A) and bad (B) quality images of the database. In both charts real images acquired with the optical, capacitive, and thermal sensor, are shown in the top row. Their respective fake images generated with cooperation are shown in the middle row, and without cooperation in the bottom row.

Clients that appear in one of the datasets (train, devel or test) do not appear in any other set. Also, for training and testing both biometric systems and anti-spoofing measures, the same indications for the usage of these subsets that appear for 2D face spoofing databases in the present document should be followed.

6.3 Protocol for Licit Biometric Transactions

In the normal operation mode, both the enrollment and the test are carried out with real fingerprints. This is used as the reference scenario. In this context the FAR (False Acceptance Rate) of the system is defined as the number of times an impostor using his own finger gains access to the system as a genuine user, which can be understood as the robustness of the system against a zero-effort attack. The same way, the FRR (False Rejection Rate) denotes the number of times a genuine user is rejected by the system.

In order to compute the performance of the system in the normal operation mode, the following sets of scores are generated: *i*) for genuine tests all the 4 real samples of each user are matched against each other avoiding symmetric matchings ($(4 \times 3)/2 = 6$ scores per user), which leads to $6 \times 24 = 144$ genuine scores, and *ii*) for impostor tests each of the four samples of every user are matched with all the samples of the remaining users in the test set avoiding symmetric matchings, resulting in $(23 \times 4 \times 4 \times 24)/2 = 4,416$ impostor scores.

6.4 Protocol for Spoofing Attacks

When considering spoofing attacks, the enrollment is achieved using real fingerprints as before, and tests are carried out with real and fake fingerprints simultaneously. In this case the genuine user enrolls with her fingerprint and the attacker tries to access the application with the corresponding gummy fingerprint. A successful attack is accomplished when the system confuses a fake fingerprint with its corresponding genuine fingerprint, i.e., the Success Rate (SR) of the attack is SR = FMR, where the FMR is the False Match Rate.

In this attacking scenario, the genuine scores are the same as computed in the protocol for licit transactions, and the impostor scores are computed matching all 4 original samples of each user with all 4 fake samples which results in $16 \times 24 = 384$ impostor scores for each scenario considered.

7 Summary

This document describes the different ICAO biometric databases of Spoofing Attacks that will be used within TABULA RASA for evaluation and development purposes. These databases are first used for the evaluation of the impact of the spoofing attacks on biometric systems and then to evaluate the effect of their respective countermeasures in the project.

References

[1] J. Fierrez, J. Ortega-Garcia, D. Torre-Toledano, and J. Gonzalez-Rodriguez. BioSec baseline corpus: A multimodal biometric database. *Pattern Recognition*, 40:1389–1392, 2007.

[2] B. Froba and A. Ernst. Face detection with the modified census transform. *IEEE International Conference on Automatic Face and Gesture Recognition*, pages 91–96, 2004.