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]

D2.5: Non-ICAO Biometric Databases of Spoofing Attacks

Due date: 29/02/2011 **Submission date:** 29/02/2011
Project start date: 01/11/2010 **Duration:** 42 months
WP Manager: Javier Acedo **Revision:** 1

Author(s): A. Riera (STARLAB), A. Soria-Frisch (STARLAB), F. Alegre (EURECOM), N. Evans (EURECOM), J. Bustard (USOU), G.-L. Marcialis (UNICA)

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

D2.5: Non-ICAO Biometric Databases of Spoofing Attacks

Abstract:

This document describes the different biometric Databases of Spoofing Attacks collected for the non-ICAO modalities. These modalities include EEG, ECG, Gait, Vein and Voice. In the case of the Vein Biometric modality, it has been decided to keep this part as confidential, and thus this section will be described in a separate document. This document also includes 3 multimodal databases descriptions: 2D-Face + Voice, 2D-Face + fingerprint and finally ECG + EEG. The different spoofing attacks are also briefly described in this document. 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 to existing biometric systems. Once the spoofing attacks have been defined (that was done in 'D2.3: Specification of Spoofing attacks'), a database containing those spoofing attacks was planned to be recorded. This deliverable describes those different databases for the non-ICAO biometric modalities, which are EEG, ECG, Gait, Vein and Voice. In the case of the Vein Biometric modality, it has been decided to keep this part as confidential, and thus this section will be described in a separate document. We have also included the description of 3 multimodal databases: 2d-Face + Voice, 2D-Face + Fingerprint and also EEG + ECG.

The structure of these deliverable is similar to its twin deliverable 'D2.4: ICAO Biometric Databases of Spoofing Attacks'. For each modality (counting the multimodal systems as individuals modalities) we have defined 4 sections: Spoofing Attacks Description (which is a summary of 'D2.3: Specification of Spoofing attacks'), Database Description and then the protocol for both the licit biometric transactions and the protocol for spoofing attacks are also described.

2 EEG Databases of Spoofing Attacks

2.1 Spoofing Attacks Description

The spoofing attack consists of a replay of the biometric signal, in this case EEG. From D2.3: The replay attack to the electro-physiological biometrics systems consists on providing a previously recorded signal to the electrodes attached to the user's body. Those electrodes tend to be quite small in order to record signals from a very specific part of the body, so the replay attack should gain access to the contact of those electrodes and replay the impostor signal through by means of an electrical contact between the electrode and the replay attack device. This electronic device shall be capable of reproducing synchronously all the previously recorded signals in all the electrodes. The impedance of the contact between this electronic device and the biometrics system electrodes might be different from the one when the electrodes are attached to the user skin. Because of that, this electronic device shall be capable of compensate the recorded signals in order to provide them as similar as they were generated by a real user.

The information needed to perform the attack is listed below:

- Biometric trait: The attacker needs to have a recording of electro-physiological signal of a subject for each of the electrodes present in the biometrics system. Those signals shall correspond to the ones from where both the electrodes and the references are placed.
- Sensor: The attackers will have more possibilities of spoofing the system if they know some key parameters of the sensor such as the sampling rate of the electrodes. The attack could fail if the provided signals were sampled in a lower sample rate than the one the sensor uses. Another important parameter of the sensor that is necessary to know in order to attack the system is the dynamic range of the analog to digital conversion the sensor performs. If the signals were over this range then the acquisition system would saturate so the recorded signal would be meaningless for the recognition system.
- Recognition Systems: No information required.

2.2 Database Description

The database has been collected in Starlab Barcelona premises. The ENOBIO sensor was used for the recording. 2 electrodes were placed in the forehead (Fp1 and Fp2 positions) for EEG, the ground in the left ear lobe and the reference in the right ear lobe. An extra electrode was also placed in the left wrist for ECG recording.

The EEG database consists of 30 subjects. Each subject has to record an enrolment set (4 2-minute takes) and a test set (15 2-minute takes).

The test set contains the countermeasures we have designed for the EEG biometric system, as explained in the next section.

2.3 Protocol for Licit Biometric Transactions

The protocol for the data recording is the following:

- For the enrolment takes the subjects will be seated and relax and they will be told not to move during the recording.

- For the test set, the subject will be in the same conditions than in the enrolment, but this time he/she will be randomly told to perform a number of actions:
  - double blinks
  - clench his teeth for a duration of about 2 seconds
  - close his/her eyes and relax for about 20 seconds
  - hold the respiration for about 20 seconds
  - contract the left hand for about 4 seconds.

The same acquisition SW is providing the instructions to the subjects, and marking the different events in the EEG signal, so we are able to know during the processing stage when each action was performed.

The first three actions (double blink, clench teeth, close eyes) were chosen because they do affect the EEG signals, and thus are useful as countermeasures. During a spoof attack, a signal generator will produce and transmit the same EEG signal to the ENOBIO sensor. If during the spoof attack, the biometric countermeasure system tells the attacker to perform a given action in a given time, the attacker will probably fail to do so, and thus would be detected as an intruder.

During this recording, ECG will also be recorded and further used as extra data.

The total amount of takes will be:

- 30 enrolments (4 2-minutes takes) = 240 minutes (EEG and ECG at the same time)
- 30*15 EEG tests (2 minutes takes) = 450 minutes (we would also have the same amount of data as extra ECG takes)

For each subject we will have 4*2+15*2=38 minutes of recording. The following list summarises all the takes recorded and the different actions performed.

- 5 takes of plain EEG (subject was not asked to perform any action)
- 10 takes in which the subject was asked to perform 2 double blinks, 2 teeth clenches and 1 close-eye task (those actions where performed in each one of the 5 takes)

In total we have 20 double blinks, 20 teeth clenches and 10 close-eyes tasks.

2.4 Protocol for Spoofing Attacks

This database will also include the replay attacks: a take recorded to subject X will be replayed by a signal generator to the ENOBIO sensor, as explained in the Spoofing Attacks Description. Each take recorded to the subjects will be used by the signal generator to perform a replay attack, so in the database we will include 30(subjects)*15(takes)=450 replay attack takes (900 minutes of EEG replay attacks).

3 ECG Databases of Spoofing Attacks

3.1 Spoofing Attacks Description

The spoofing attack consists of a replay of the biometric signal, in this case ECG. Please refer to section 2.1, since the spoofing attack here is identical to the one described above.

3.2 Database Description

The data base has been collected at the same time than the EEG one. Please refer to section 2.3 for further details.

3.3 Protocol for Licit Biometric Transactions

The protocol for the data recording is the following:

- For the enrolment takes the subjects will be seated and relax and they will be told not to move during the recoding.

- For the test set, the subject will be in the same conditions than in the enrolment, but this time he/she will be randomly told to perform a number of actions:
  - double blinks
  - clench his teeth for a duration of about 2 seconds
  - close his/her eyes and relax for about 20 seconds
  - hold the respiration for about 20 seconds
  - contract the left hand for about 4 seconds.

The same acquisition SW is providing the instructions to the subjects, and marking the different events in the ECG signal, so we are able to know during the processing stage when each actions was performed.

The last two actions (hold breath and contract left hand) were chosen because they do affect the ECG signals, and thus are useful as countermeasures. During a spoof attack, a signal generator will produce and transmit the same ECG signal to the ENOBIO sensor. If during the spoof attack, the biometric countermeasure system tells the attacker to perform a given action in a given time, the attacker will probably fail to do so, and thus would be detected as an intruder.

During this recording, EEG will also be recorded and further used as extra data.

The total amount of takes would be:

- 30 enrolments (4 2-minutes takes) = 240 minutes (EEG and ECG at the same time)

- 30*15 EEG tests (2 minutes takes) = 450 minutes (we would also have the same amount of data as extra EEG takes)

For each subject we would have 4*2+15*2=38 minutes of recording. The following list summarises all the takes recorded and the different actions performed.

- 5 takes of plain ECG (subject was not asked to perform any action)
- 5 takes in which the subject was asked to perform 4 left hand contractions, 1 hold respiration task (those actions where performed in each one of the 5 takes)
- 5 takes in which the subject was asked to perform 2 left hand contractions, 1 hold respiration task (those actions where performed in each one of the 5 takes)

In total we have 30 left hand contractions and 10 hold respiration tasks.

3.4 Protocol for Spoofing Attacks

The protocol of the spoofing attack is the same as the EEG one. Please refer to section 2.4 for further details.

4 Gait Databases of Spoofing Attacks

4.1 Spoofing Attacks Description

There is currently one commercial gait recognition systems in development that will soon be released by a consortium including the National Physical Laboratory. However as this system is not currently available for analysis the Southampton Biometric Tunnel is being used as an alternative. This is particularly appropriate as both systems use a silhouette based gait recognition technique on which no existing spoofing attacks have been attempted. As mentioned in deliverable 2.3, simple replay attacks of gait are unlikely to be feasible due to the practical difficulties of simultaniously spoofing multiple calibrated and synchronised cameras. It is therefore necesary to invent a novel gait spoofing technique. To achieve this, the factors that contribute significantly to the gait signature need to be investigated. Of particular importance are the factors that can be easily and consistently altered to spoof identity. In particular, silhouette based gait systems are potentially vulnerable to spoofing attacks where an attacker replicates the silhouette of the target. The most straightforward and unobtrusive method for performing this style of attack is to wear clothing that makes an attacker's bodyshape appear the same as the target. This can be viewed as a similar to wearing a 3D mask in order to spoof a facial recognition system.

One additional consideration is that 3D gait recognition systems can observe the full volume of the subject, making it relatively impractical for large subjects to spoof smaller subjects as they cannot add anything to reduce their volume. Therefore it should be expected that any given subject will only be able to spoof some of the other subjects.

As this is the first clothing-based spoofing attack the investigation focuses on the simplest form of attack, which is to replicate the clothing of the target. This is an important potential vulnerability as such an attack is relatively straightforward. It is also likely that such an approach will be already used by an attacker to unobtrusively enter a secure area where uniforms or formalised styles of dress are common.

4.2 Database Description

The database consists of 22 subjects (14 male and 8 female), between 20-55 years old. The subjects were recorded walking through the Southampton Gait Tunnel in both their normal clothes and whilst wearing a common uniform. By having every subject wear the same clothes the degree to which one subject could impersonate another by mimicing their clothes can be investigated. The uniform clothing appearance was achieved by having subjects wear white overalls over their normal clothes. Each recording of normal or uniform clothing was repeated between 10 and 35 times depending on subject availability. The only instruction to subjects was to walk normally. The gait tunnel ensures that background, lighting, walking surface, and calibrated position of cameras remain constant across all recordings.

4.3 Protocol for Licit Biometric Transactions

Licit biometric transactions are straightforward, the subject walks naturally through a corridor with calibrated cameras and a known background, in this case the Southampton Biometric Tunnel. In the use case of covert identity verification (usecase 2.5 in described in D2.1) the resulting signature is compared against a database of known subjects. In the case of Physical Access Control (usecase 2.2 as described in D2.1) the subject uses some form of claimed identity, such as a keycard, which is then used to verify that the recorded gait signature matches the enrolled gait signature.

The database is recorded using the same camera, lighting and other factors as the baseline database and can therefore be used as an extension. This makes it possible to use the larger baseline database to provide a relatively accurate false reject rate. The baseline then provides thresholds which can be used to calculate the probability of a spoofing success, measured as the false accept rate under spoofing attack.

4.4 Protocol for Spoofing Attacks

All subjects participated in spoofing attacks and the procedure did not differ from a licit transaction, other than that the attackers were wearing the same clothing as the target. The database also contains recordings of subjects wearing normal clothes. These recordings are used to evaluate the importance of matching the build of the target in addition to matching their clothing.

Deliverable 3.4 describes the evaluation of targeted clothing, where an attacker selects a target with a similar build and impersonates their clothing, it includes three metrics: Targeted attack false accept rate Clothing attack false accept rate Targeted clothing attack false accept rate

Targeted attacks are measured by comparing each of the normal clothes recordings of each subject against each of the normal clothes recordings of the subject with most similar build, estimated by selecting the subject whose signature most is falsely accepted most frequently at the equal error rate threshold of the baseline. In total this represents 21,392 measurements. Clothings attacks are calculated by comparing each of the uniform recordings of each subject against the uniform recordings of all of the other subjects. In total this represents 2,268,642 measurements. Targeted clothing attacks are the same as targeted attacks except instead of using the normal clothing recordings the uniform clothing recordings are used. This is equivalent to each subject selecing the person with the most similar build and impersonating their clothing. In total this database can be used to calculate 92,870 measurements.

The number of recordings per person and the relative number of clothing to normal recordings varies with subject due to their differing availability for recording sessions. However all subjects have a minimum of 10 recordings wearing normal clothes and 49 wearing a common uniform.

5 Voice Databases of Spoofing Attacks

5.1 Spoofing Attacks Description

As per previously reported in D2.3, for the voice biometric we will investigate the following attacks:

1. **Replay attacks:** as a form of baseline spoofing which aims to assess the effect of using different transducers to capture the speech signal and replaying it to the biometric system.

2. **Artificial signals:** which can be used on their own or injected into an attackers natural voice signal in order to boost the system likelihood or score.

3. **Voice conversion:** aims to transform an attacker's voice toward that of a client. We will concentrate mostly on this form of attack.

4. **Voice synthesis:** will only be addressed as a proof of concept and to compare its efficacy with other attacks.

The four attacks will be investigated using two databases. The first corresponds to a large-scale, though somewhat artificial sensor-level attack setup. The second relates to a smaller-scale, though more practical setup which is entirely consistent with sensor-level attacks. The two setups also reflect the voice modality use case scenarios reported in D2.1. They are the physical and mobile access scenarios respectively. The two setups are described in more detail below.

5.2 Database Description

Spoofing databases will be created according to one of two scenarios: mobile/telephony and physical-access.

Mobile/telephony (NIST baseline):
The mobile/telephony scenario will be addressed in TABULA RASA using NIST Speaker Recognition Datasets[1], namely the exact same datasets that were used for baseline evaluations reported in D3.2. In summary, the NIST'04 datasets are used for background, normalisation or session modelling, the NIST'05 datasets are used for development and the NIST'06 dataset is used for evaluation. Further details regarding the specification of each dataset are freely available from NIST's website. All data used to effect spoofing attacks will come from the same datasets.

The only difference between the data used for baseline experiments, reported in D3.2, and the data to be used for spoofing relates to the use of a different experimental condition.

[1] http://www.itl.nist.gov/iad/mig/tests/spk/

Instead of the 1conv4w condition used for baseline experiments, spoofing and related experiments will be performed using data from the 8conv4w condition which provides multiple sessions for each speaker. It contains voice recordings from 502 speakers (211 male, 291 female). There are 8 sessions for each speaker giving a total of $502 \times 8 = 4016$ recordings of approximately 5 minutes in duration. Those sessions not used for testing are used to effect spoofing attacks for all impostor trials in the standard NIST protocols.

In addition to the new baseline dataset (without spoofing), three new datasets have been generated where all impostor test segments are replaced with spoofed versions coming from artificial signals, voice conversion or voice synthesis. Replay attacks are not assessed in the case of mobile/telephony data since the somewhat artificial nature of the sensor-level attacks to be investigated with NIST SRE datasets means that replay attacks will be impossible to detect without using challenge-response countermeasures; they will not differ from a genuine client trial.

Under the assumption that all system elements can be modelled as different, linear processes (as is the case for the three attacks investigated), then spoofing, acquisition and transmission (channel) is equivalent to acquisition and transmission followed by spoofing. Under the assumption of linearity, spoofing work involving NIST SRE datasets can therefore be performed directly on the exact same, pre-recorded datasets, without re-recording, while still being consistent with sensor-level spoofing.

Physical access (EURECOM database, MOBIO baseline):
To support the physical access scenario and to evaluate spoofing through an approach entirely consistent with sensor-level attacks, we have also collected a new database with which to assess spoofing. Except for the use of a different sensor, its specification is similar to that of the MOBIO database for which baseline results were reported in D3.2. However, it was not possible to collect sufficient, independent data to learn new background models, specific to the different sensor and recording conditions and thus all work conducted with the new 'EURECOM' database will use the background model learned on MOBIO data. As a consequence of using different sensors and recording conditions, recognition results on the new dataset are worse than those for the baseline MOBIO database. The effect of background mis-match is however the same for genuine client trials and spoofed impostor trials and thus the setup is nonetheless valid for the assessment of spoofing.

The EURECOM dataset consists of voice recordings from 21 speakers (14 male, 7 female) who each provide recordings in 4 sessions. All data is captured with a Logitech webcam and, as per the original MOBIO dataset, each session involves the recording of 5 short phrases (short questions), 3 read sentences and 7 longer, free-text phrases (longer questions). The database contains a total number of $21 \times 15 \times 2 = 630$ recordings. Given the modest size of the database, all data used to effect spoofing comes from within the test dataset itself. Since this always comes from a different, target person (and not the impostor), this does not present any issues regarding data independency.

In addition to a new baseline dataset (without spoofing), four different datasets are generated for each of the four spoofing attacks: replayed data, artificial signals, voice conversion and voice synthesis.

Mobile/telephony scenario - NIST SRE datasets			
		Development (NIST'05)	Evaluation (NIST'06)
Baseline	Male Female	211/1231/10608 291/1337/10726	251/1543/14297 343/1843/17239
Replay Attack	Male Female	-	
Artificial Signals	Male Female	211/1231/1231 291/1337/1337	251/1543/1543 343/1843/1843
Voice Conversion	Male Female	As for baseline	
Speech synthesys	Male Female	As for artificial signals	

Table 1: An illustration of the size of each dataset used for spoofing assessment in the case of the mobile/telephony scenario (NIST datasets). Figures illustrate the number of clients / genuine client trials / and impostor trials.

5.3 Protocol for Licit Biometric Transactions

We describe here protocols for the two scenarios and corresponding datasets to be used for all future voice recognition work in TABULA RASA.

Mobile/telephony (NIST baseline):
Except for the use of the 8conv4w condition instead of the 1conv4w condition, protocols are exactly the same as those for all baseline results reported in D3.2. Once again, the NIST'04 dataset is used for background data, the NIST'05 dataset is used for development and the NIST'06 dataset is used for evaluation. All experiments relate to the core condition which involves approximately 5 minutes of training and testing data.

The new protocols/conditions result in a slightly reduced number of speakers and recordings than used for previously reported baseline experiments. A summary of the dataset sizes illustrating the number of clients, genuine client trials and impostor trials is illustrated in Table 1 for both development and evaluation datasets. The development set contains 502 clients (211 male, 291 female) whereas the evaluation set contains 594 clients (251 male, 343 female). These numbers are consistent for the baseline protocol and the three spoofing protocols (described below). The precise protocol is non-exhaustive and exactly as defined by NIST. They give 2568 and 21334 genuine client and impostor trials respectively for the development set and 3386 and 31536 genuine client and impostor trials respectively for the evaluation dataset.

Physical access (EURECOM database, MOBIO baseline):
Except for there being fewer clients and sessions the protocol for the EURECOM database is similar to that of the MOBIO database. A summary of the dataset sizes illustrating the number of clients, genuine client trials and impostor trials is illustrated in Table 2 for

Physical access scenario - EURECOM datasets			
		Development	Evaluation
Baseline	Male Female	7/14/84 3/6/12	7/14/84 4/8/24
Replay Attacks	Male Female	7/14/14 3/6/6	7/14/14 4/8/8
Artificial Signals	Male Female	As for replay attacks	
Voice Conversion	Male Female	As for baseline	
Speech synthesys	Male Female	As for replay attacks	

Table 2: An illustration of the size of each dataset used for spoofing assessment in the case of the physical access scenario (EURECOM datasets). Figures illustrate the number of clients / genuine client trials / and impostor trials.

both development and evaluation datasets. The development set contains 10 clients (7 male, 3 female) and the evaluation set also contains 10 clients (7 male, 4 female). These numbers are consistent for the baseline protocol and the three spoofing protocols (described below). Only data from the first session and short answers (questions 1 to 5) is used for training or enrollment whereas data from the second, third and fourth sessions and read text (questions 6 to 8) and free text (questions 9 to 15) are used for testing. Scores are generated in an exhaustive fashion thereby using every possible combination of target and impostor tests. The number of genuine client and impostor trials is 20 for the development set and 22 for the evaluation datasets respectively.

We fully acknowledge that the number of persons and trials is small and will not permit any considerable statistical significance in generated results. Cross-validation will be performed (e.g. Session 4 for training and Sessions 1-3 for testing) and, although this will help to increase the number of tests, results will nonetheless relate to a small number of clients. In this sense it will be of interest to compare results for the small, but practical setup used for the physical access scenario to those for the larger, though somewhat artificial setup used for the mobile/telephony scenario.

5.4 Protocol for Spoofing Attacks

The protocols for spoofing assessment are identical to the licit protocol only that, for each impostor access attempt, the test sample is treated or replaced according to the given spoofing attack. The protocols for the NIST SRE and EURECOM datasets differ slightly and are described below.

Mobile/telephony (NIST baseline):
For the NIST SRE datasets, any data used to effect spoofing comes from one of the sessions not used for testing. Of the 8 available, the first is used for testing, thus sessions 2 to 8 can be used for spoofing purposes. Furthermore, any other suitable data, e.g. the NIST 2008 dataset, will be used as independent background data (i.e. for learning a universal background model used in voice conversion). Except for modifications to impostor test segments through spoofing, protocols are exactly the same as described above. Accordingly there is no overlap between the data used to train client models and that used for spoofing.

The number of genuine client trials remains the same as described above, however, the number of impostor trials is dependent on the form of spoofing attack. As illustrated in Table 1 the number of impostor trials for artificial signals is greatly reduced (1231 c.f. 10608 for male and 1337 c.f. 10726 for female). This is because, when comparing a voice model for person A with an impostor test segment B, the impostor test segment is replaced with an artificial signal which is targetted only towards model A and is independent to impostor B. This is also the case for voice synthesis. With voice conversion, however, a transformation is learned which maps the impostor test segment B toward model A. It is specific to the test segment and thus the number of impostor tests in this case is the same as that in the baseline. As a result, with artificial signals and speech synthesis spoofing attacks the number of impostor trials is reduced to 1231 male and 1337 female tests for the development set and 1543 male and 1843 female tests for the evaluation set.

Physical access (EURECOM database, MOBIO baseline):
As in mobile/telephony scenario, protocols for spoofing attacks are dependent on the spoofing attack which are again effected by modifying or replacing impostor test data according to any of the four spoofing attacks. For the same reasons as outlined above and as illustrated in Table 2, for replay attacks, artificial signals and speech synthesis, the number of impostor trials is reduced to 14 male and 6 female tests for the development set and to 14 male and 8 female for the evaluation set. For voice conversion the number of trials is identical to that in the baseline. In this case, due to the small dataset size any data used for spoofing may come from the same session used for testing (but different speaker), but never from the same session as used for training client models. Since only impostor test segments are modified or replaced, this does not present any issues with regard to independency. Note also that, given the lack of data, the same world model used for licit transactions may be used for spoofing.

6 Multi-modal Databases of Spoofing Attacks: 2D-Face and Voice

6.1 Spoofing Attacks Description

Attacks of voice and 2d faces against multi-modal systems are basically the same considered for the face-fingerprint multi-modal system:

- Direct attack to the voice matcher. In this case, only voice biometric is spoofed, whilst the input to the face matcher is, typically, an impostor trial.
- Direct attack to the face matcher. In this case, only face biometric is spoofed, whilst the input to the voice matcher is, typically, an impostor trial.
- Direct attack to both matchers. In this case, the attacker has been able to replicate both biometrics of the targeted client.

In multi-modal systems, the consequence is that fused match scores will be a sort of "average" between a standard impostor attack and a genuine user trial. The worst-case is that both biometrics have been spoofed, which makes the response of the system similar to that of a mono-modal one.

6.2 Database Description

In order to simulate above attacks, we may consider that data set is characterized by "true" attacks, that is, captured images are from real attacks (for example photos in front of the webcam, or a counterfeited/imitated voice).

This data set is made up of a set of genuine faces and fingerprints, and of a set of replicated faces and voices, corresponding to the same users of the data set.

6.3 Protocol for Licit Biometric Transactions

Given the following assumptions:

- let C be the number of clients;
- let N be the Number of voice/face samples for each client;
- the same number of voice and fingerprint images is assumed to be taken for each client;
- let N_{dev} and N_{test} be the number of voice/face samples considered for the development (namely, templates and fusion rules parameters setting) and for testing the system, respectively. Obviously $N = N_{dev} + N_{test}$.

The protocol for licit biometric transactions is as follows:

1. For i = 1...C
   (a) For j = 1..N_{test}
      i. Consider takes j of client i, one as voice and one as face.
      ii. Above takes are compared with the templates of client i, for both matchers M_{voice} and M_{face};
      iii. Systems outputs are match scores $s_{j,voice}$ and $s_{j,face}$;
      iv. A fusion rule $f()$ is applied in order to obtain $s_j = f(s_{j,voice}, s_{j,face}$.
      v. If $s_j > s^*$, where s^* is the operational threshold, the claimed identity is verified, and the related client is classified as client i.

The aim of this protocol is to obtain the testing genuine users distributions, which will be compared with the attack distributions computed as described in the next Section.

6.4 Protocol for Spoofing Attacks

We consider only the case in which we have realistic spoofing attacks. Therefore, we assume that:

- let C be the number of clients;
- let N_{test} be the number of fake voice/face samples considered for testing the system, respectively.

The protocol for spoofing attacks is as follows:

1. For i = 1...C
   (a) For j = 1..N_{test}
      i. Consider fakes j of client i, one as voice and one as face.
      ii. Above takes are compared with the templates of client i, for both matchers M_{fing} and M_{face};
      iii. Systems outputs are match scores $s_{j,voice}$ and $s_{j,face}$;
      iv. A fusion rule $f()$ is applied in order to obtain $s_j = f(s_{j,voice}, s_{j,face}$.
      v. If $s_j > s^*$, where s^* is the operational threshold, the claimed identity is verified, and the related client is classified as client i.

Above protocol is aimed to describe multi-modal spoofing attacks. In the case of mono-modal attacks, e.g. when only voices are spoofed, the corresponding match score will be coupled with an impostor face match score randomly chosen among the other clients.

The aim of this protocol is to obtain the testing fake distributions, which will be compared with the genuine users distributions computed as described in the previous Section.

7 Multi-modal Databases of Spoofing Attacks: 2D-Face and Fingerprint

7.1 Spoofing Attacks Description

The attacks to multi-modal biometric systems based on faces and fingerprints can be summarized as follows:

- Direct attack to the fingerprint matcher. In this case, only fingerprint biometric is spoofed, whilst the input to the face matcher is, typically, an impostor trial.
- Direct attack to the face matcher. In this case, only face biometric is spoofed, whilst the input to the fingerprint matcher is, typically, an impostor trial.
- Direct attack to both matchers. In this case, the attacker has been able to replicate both biometrics of the targeted client.

In multi-modal systems, the consequence is that fused match scores will be a sort of "average" between a standard impostor attack and a genuine user trial. The worst-case is that both biometrics have been spoofed, which makes the response of the system similar to that of a mono-modal one.

7.2 Database Description

In order to simulate above attacks, we may consider that data set is characterized by "true" attacks, that is, captured images are from real attacks (for example photos in front of the webcam, or a fake fingerprint).

This data set is made up of a set of genuine faces and fingerprints, and of a set of replicated faces and fingerprints, corresponding to the same users of the data set.

7.3 Protocol for Licit Biometric Transactions

Given the following assumptions:

- let C be the number of clients;
- let N be the Number of fingerprint/face images for each client;
- the same number of face and fingerprint images is assumed to be taken for each client;
- let N_{dev} and N_{test} be the number of fingerprint/face images considered for the development (namely, templates and fusion rules parameters setting) and for testing the system, respectively. Obviously $N = N_{dev} + N_{test}$.

The protocol for licit biometric transactions is as follows:

1. For i = 1...C
   (a) For j = 1..N_{test}
      i. Consider takes j of client i, one as fingerprint and one as face.
      ii. Above takes are compared with the templates of client i, for both matchers M_{fing} and M_{face};
      iii. Systems outputs are match scores $s_{j,fing}$ and $s_{j,face}$;
      iv. A fusion rule $f()$ is applied in order to obtain $s_j = f(s_{j,fing}, s_{j,face}$.
      v. If $s_j > s^*$, where s^* is the operational threshold, the claimed identity is verified, and the related client is classified as client i.

The aim of this protocol is to obtain the testing genuine users distributions, which will be compared with the attack distributions computed as described in the next Section.

7.4 Protocol for Spoofing Attacks

We consider only the case in which we have realistic spoofing attacks. Therefore, we assume that:

- let C be the number of clients;
- let N_{test} be the number of fake fingerprint/face images considered for testing the system, respectively.

The protocol for spoofing attacks is as follows:

1. For i = 1...C
   (a) For j = 1..N_{test}
      i. Consider fakes j of client i, one as fingerprint and one as face.
      ii. Above takes are compared with the templates of client i, for both matchers M_{fing} and M_{face};
      iii. Systems outputs are match scores $s_{j,fing}$ and $s_{j,face}$;
      iv. A fusion rule $f()$ is applied in order to obtain $s_j = f(s_{j,fing}, s_{j,face}$.
      v. If $s_j > s^*$, where s^* is the operational threshold, the claimed identity is verified, and the related client is classified as client i.

Above protocol is aimed to describe multi-modal spoofing attacks. In the case of mono-modal attacks, e.g. when only fingerprints are spoofed, the corresponding match score will be coupled with an impostor face match score randomly chosen among the other clients.

The aim of this protocol is to obtain the testing fake distributions, which will be compared with the genuine users distributions computed as described in the previous Section.

8 Multi-modal Databases of Spoofing Attacks: ECG and EEG

8.1 Spoofing Attacks Description

The spoofing attacks to be collected for the multimodal EEG-ECG attacks will be included in 3 different databases. The general procedure for the generation of the spoofing attacks has been described in Sec. 2.1. Basically the spoofing attacks are generated by a device that replays the physiological signals used for authentication. The reader is referred to that section for further details.

8.2 Database Description

The performance evaluation of spoofing attacks for the multi-modal authentication based on EEG and ECG will be realized with 3 different databases. Analogously to other multimodal attacks these 3 databases correspond to 3 different attack scenarios (e.g. see Sec. 6.1):

1. Direct attack to the EEG modality. In this case, only the EEG biometric is spoofed, i.e. reproduced with the aforementioned device, whereas the input to the ECG modality is a legal trial (possibly stolen from the database).

2. Direct attack to the ECG modality. In this case, only the ECG biometric is spoofed, i.e. reproduced with the aforementioned device, whereas the input to the EEG modality is a legal trial (possibly stolen from the database).

3. Direct attack to both modalities. In this case, the attacker has been able to replicate, i.e. reproduce with the aforementioned device, both biometrics of the target client.

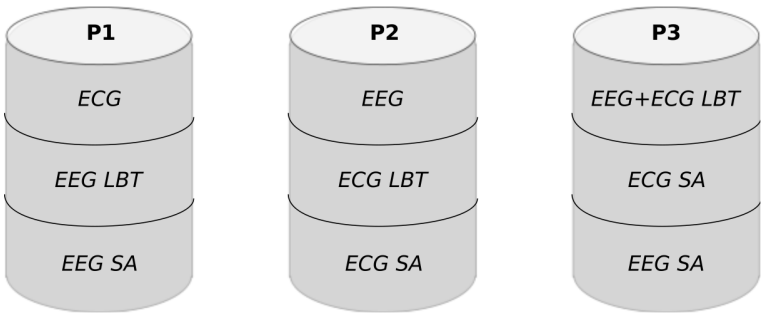

Figure 1: Three different databases acquired following the protocols for EEG attacks (P1), ECG attacks (P2), and simultaneous EEG-ECG attacks (P3). Each of these databases includes the data subsets acquired with the corresponding protocols for licit biometric transactions (LBT), and spoofing attacks (SA).

Based on these scenarios 3 different databases can be generated with different data acquisition protocols, which have been denoted as P1, P2, and P3. A scheme of databases' content is depicted in Fig.1. Each of these protocols include a protocol for licit biometric transactions (LBT) and for spoofing attacks (SA) as described in the following sections.

8.3 Protocol for Licit Biometric Transactions

The protocol for Licit Biometric Transactions of P1, and P2, have been described respectively in Secs. 2.3, and 3.3. The data acquired in those unimodal sessions, which are denoted as EEG LBT and ECG LBT in Fig. 1, will be used as well for the evaluation of the multimodal setting. Besides these data subsets, the data acquisition protocol P3, which takes into consideration a scenario for simultaneous attack of the EEG and the ECG modalities, will be used for generation the corresponding data subset, i.e. EEG+ECG LBT in P3 (see Fig. 1). P3 is described in the following paragraphs.

The ENOBIO sensor will be used for the recording. Two electrodes will be placed in the forehead (Fp1 and Fp2 positions) for EEG, the ground in the left ear lobe, and the reference in the right ear lobe. An extra electrode will be placed also in the left wrist for ECG recording. The EEG+ECG LBT data subset will correspond to data of 30 subjects. These subjects will be first enrolled in the system. This procedure will take place during the data acquisition following the protocol described in Sec. 2.3, so we do not need to record the enrollment data again here. These enrollment data is acquired while the subject seated, relaxed, and without movement. The enrollment data for each subject accounts for 4 times 2-minutes takes, i.e. for a total of 240 minutes for the 30 subjects (simultaneous EEG and ECG acquisition).

In the protocol for the acquisition of the EEG+ECG LBT data subset, the subjects will be asked to perform some actions. These actions are expected to be evaluated as counter-measures in further stages of the project. The list of actions include the following:

- In the first 40 seconds she will be asked to perform 3 double blinks and 3 triple blinks.
- In the next 40 seconds she will be asked to clench her teeth 6 times for a duration of about 2 seconds.
- In the next 40 seconds she will be asked to keep her eyes closed and relax for the last 20 seconds.
- In the next 60 seconds she will be asked to perform 10 hand contractions (open and close the left hand).
- In the last 60 seconds interval, particularly from second 60 to second 80, she will be asked to hold her breath for about 20 seconds.

The total amount of data in this subset will be: 30 subjects times 5 EEG+ECG sessions times 4 minutes takes, which account for up a total data amount of 600 minutes. For each subject we will have 5*4=20 minutes of EEG+ECG LBT recording.

8.4 Protocol for Spoofing Attacks

The data sets generated following the three protocols P1, P2, and P3 (see Fig. 1) will include some spoofing attack data. The data subsets corresponding to the spoofing attacks will be generated in the protocols described in Secs. 2.4, and 2.4. These data subsets have been denoted respectively as EEG SA and ECG SA. As already mentioned the data included in these subsets result from replaying the corresponding electrophysiological signals with a suitable device. Each data subset includes 30*5=150 replay attack takes (300 minutes of replay attacks). While EEG SA is used in the databases acquired following P1 and P3, the ECG SA subset is included in the databases corresponding to P2 and P3 (see Fig. 1).

9 Summary

This document describes the different Non-ICAO biometric databases of Spoofing Attacks that will be used within TABULA RASA for evaluation purposes. The section describing the Vein modality will be delivered in a separate document, since it has been decided to keep that section confidential. These databases will be used for the evaluation of the success of the spoofing attacks and also to evaluate the effect of their respective countermeasures later on in the project.