BEAT

Biometrics Evaluation and Testing

http://www.beat-eu.org/

Funded under the 7th FP (Seventh Framework Programme)

Theme SEC-2011.5.1-1
[Evaluation of identification technologies, including Biometrics]

D3.1: Description of Initial Reference Biometric Systems

Due date: 26/06/2012 **Submission date:** 25/05/2012
Project start date: 01/03/2012 **Duration:** 48 months
WP Manager: Julian Fierrez **Revision:** 1

Author(s): J. Galbally (UAM), J. Fierrez (UAM), A. Anjos (IDIAP), S. Marcel (IDIAP), N. Poh (UNIS), J. Kittler (UNIS), C. C. Ho (UNIS), J. Bringer (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

D3.1: Description of Initial Reference Biometric Systems

Abstract:

This document defines the different initial reference biometric systems that will be considered within the BEAT project. The range of biometrics considered includes: 2D face, fingerprint, iris and vein, in addition to multi-modal biometrics. The systems described here will be used for baseline evaluations within the project. Thus, formal specifications should be given.

Contents

1 Introduction

This document details the initial reference biometric systems that will be considered within the European Union (EU) 7th Framework Programme (FP7) Small or Medium-Scale Focused Research Project (STREP) entitled 'Biometrics Evaluation and Testing' (BEAT).

Biometrics included within the focus of the project include: 2D face, fingerprint, iris, and vein. The project will also consider various multi-modal biometric combinations. In most cases there is more than one system described per biometric, mostly due to the need for compatibility with multi-modal biometrics for which performance will be compared to their respective mono-modal counterparts.

The description of several systems for each modality will permit a deeper understanding of the different traits and a wider range of results in the evaluations carried out throughout the project. However, not every system will be necessarily used in every assessment experiment. Which of them will be selected for the different tests will largely depend on the experiment at hand. It is therefore very important that the context of each evaluation is clearly defined.

This document does not describe any evaluation work, baseline or otherwise. This will be reported in following deliverables. Evaluation work, both of baseline systems performance and security threats and countermeasures will, however, be based in general upon the reference systems described in this document.

The rest of this document is organised as follows. Each mono-modal biometric is first discussed in Sections 2 to 5. In each case the selected systems are described with a common structure. Finally, multi-modal systems are described in Section 6.

2 Face biometrics

Face recognition aims to uniquely recognize individuals based on their facial physiological attributes. It is a very active field of research because face recognition is natural and non-intrusive. Despite the significant progress in the field, the technology does not yet meet all security and robustness requirements needed by an authentication system for unsupervised deployment in high security environments. In addition to difficulties related to robustness against a wide range of viewpoints, occlusions, ageing of subjects and complex outdoor lighting, face biometric techniques have also been shown to be particularly vulnerable to various kinds of attacks [1].

2.1 Existing systems

Face recognition as a research field has existed for more than 30 years and has been particularly active since the early 1990s [16]. Researchers from many different fields (from psychology, pattern recognition, neuroscience, computer graphic and computer vision) have attempted to create and understand face recognition systems [16]. This has led to many different methods and techniques in this field.

These techniques have often been divided into two groups (1) holistic matching methods and (2) feature-based matching methods [16]. Holistic approaches use the whole face as one input while feature-based methods extract multiple features (eye position, nose position, angles between physical features or local frequency responses) and analyze them separately before fusing the different results [16]. However, recently several of the most advanced methods can be considered both feature-based (part-based) and holistic, therefore, in this document it is no longer considered useful to make a strict division between holistic and feature-based methods.

The state-of-the-art for face recognition is currently dominated by two themes: using parts or partitions of the face (which is often not strictly a feature-based nor a holistic technique) and the use of Local Binary Patterns (LBPs) [14]. These two themes are exemplified by the local Gabor binary histogram sequences (LGBPHS) technique [18]. This technique obtains local histograms of LBPs from non-overlapping blocks and then concatenates these histograms to form a single sequence or feature vector; this can be considered to be both feature-based and holistic. These two themes are not unique to the LGBPHS technique, with several other methods making use of them as the basis for their systems, this includes: region-based LBP histograms [20] with adaptation [17] and feature distribution modeling of the local discrete Cosine transform (DCT) [12]. In the following paragraphs we describe in more detail some of the state-of-the-art face recognition systems.

LBP histograms

In 2004 Ahonen et al. [20] proposed a face recognition system using LBP histograms. The technique first applied the LBP to each pixel of the face image and then divided this encoded face image into a set of regions. Histograms were then obtained from each region and then concatenated to form a single feature vector. This work has been

used and extended by several authors to incorporate adaptation of the histograms using a background model [17], reduce the size of each histogram by performing linear discriminant analysis (LDA) [15] and to replace histogram counts with a kernel density estimation method [19].

Local Gabor binary pattern histogram sequences

The *Local Gabor Binary pattern Histogram Sequences (LGBPHS)* technique [18] applies Gabor wavelets at multiple scales and orientations to obtain several sub- images. These sub-images are then encoded using a standard LBP operator and these local Gabor binary maps are then divided into non-overlapping regions (parts) and a histogram is computed on each region (part). The histograms, for each region for each map, are compared using the histogram intersection measure.

Feature distribution modeling of face parts

Feature distribution modeling of face parts has taken two approaches, to consider the face as a set of: independent observations and dependent observations. The first method, the GMM part-based approach [11], divides the face into blocks obtains DCT features from each block and models these features using GMMs; it assumes that each block (or feature) is an independent observation of the face. The second method, the HMM part-based approach [13], divides the face into parts and then models the feature vectors either down the face using a single HMM or down the face and then across the face using a main HMM with embedded HMMs. For both techniques it has been shown to be advantageous to use the LBP as a pre-processing technique [12].

There are other techniques which consider face recognition is a considerably difference way. One example is the tied-factor analysis method described by Prince et al. [9] for dealing with pose variation. Also, the recent technique proposed in [10] where the face is described by expert classifiers for features such as hair, skin color, and overall appearance provides a completely different approach; one caveat with this technique is that these classifiers must be derived by first annotating a database with all of these features with several human operators.

2.1.1 System 1: Part-Based GMM System

A part-based gaussian mixture model (PB-GMM) system is chosen as baseline for face authentication. The system combines a part-based face representation and GMMs. It divides the face into blocks, and treats each block as a separate observation of the same underlying signal (the face). A feature vector is thus obtained from each block by applying the Discrete Cosine Transform (DCT) [7]. The distribution of the feature vectors is then modeled using GMMs.

For feature extraction, the face is normalized, registered and cropped. This cropped and normalized face is divided into blocks (parts) and from each block (part) a feature vector is obtained. Each feature vector is treated as a separate observation of the same underlying

signal (in this case the face) and the distribution of the feature vectors is modeled using GMMs. The feature vectors from each block are obtained by applying the DCT.

Once the feature vectors are calculated, feature distribution modelling is achieved by performing background model adaptation of GMMs [3, 4]. Background model adaptation first involves the training a world (background) model Ω_{world} from a set of faces and then the derivation of client models Ω^i_{client} for client i by adapting the world model to match the observations of the client. The adaptation is performed using a technique called mean only adaptation [5].

To verify an observation, x, it is scored against both the client (Ω^i_{client}) and world (Ω_{model}) model. The two models, Ω^i_{client} and Ω_{world}, produce a log-likelihood score which is then combined using the log-likelihood ratio (LLR) to produce a single score. This score is used to assign the observation to the world class of faces (not the client) or the client class of faces (it is the client) based on a predefined threshold τ.

The implementation of choice for Face Recognition will be based on Bob [8, 21], a freely and publicly available framework for signal processing and machine learning. The framework can be deployed for out-of-the-box face processing and includes face detection, normalization, feature extraction, model training and matching.

3 Fingerprint biometrics

Fingerprint biometrics is one of the most developed biometric technologies, with multiple commercial products and adequate performance levels for applications such as physical access control, given that the population considered and the acquisition scenario are controlled and well behaved. These solutions can be nevertheless inefficient or even impracticable when confronted with the varying quality of data encountered in some applications such as forensics in realistic scenarios, where latent fingerprints with low quality and partial data are usually encountered. The main focus of the BEAT project regarding fingerprint biometrics is on evaluating state-of-the-art commercial systems on controlled data, evaluating its vulnerabilities, and finally developing adequate countermeasures against those vulnerabilities. The application of automated fingerprint biometrics in forensics with latent data is out of the scope of the project.

3.1 Systems

As the market leading biometric many different fingerprint recognition systems have been proposed in the literature. From a general point of view all of them may be included in one of these three categories: *i*) correlation-based, *ii*) minutiae-based, or *iii*) based on features of the ridge pattern.

3.1.1 Existing systems

Correlation-based methods
These systems compute and maximize the cross correlation between pixels of the stored and input samples. As an effect of the displacement and the rotation that exists between samples of the same fingerprint, the similarity cannot be computed by simply evaluating the correlation but it has to be maximized for different vertical and horizontal offsets of the fingerprint, and for different rotations. These operations entail a huge computational cost and, except for very good quality samples, these methods do not present in general comparable results to those obtained with the other two types of approaches. Different algorithms have been developed that are able to substantially decrease the computational cost of the matching process by computing the local correlation at specific areas of the fingerprint such as the core, or close to very good quality minutiae points. However, none of these techniques provide a clear performance improvement over the general method.

Minutiae-based methods
This is the most popular and widely used technique as it presents the best performance results, and is the basis of the fingerprint comparison made by fingerprint examiners. Minutiae are extracted from the two fingerprints and stored as sets of points in the two-dimensional plane. Minutiae-based matching essentially consists of finding the alignment between the template and the input minutiae sets that result in the maximum number of minutiae pairings.

Ridge feature-based methods
Minutiae extraction is difficult in very low-quality fingerprint images. However, whereas other features of the fingerprint ridge pattern (e.g. local orientation and frequency, ridge shape, texture information) may be extracted more reliably than minutiae, their distinctiveness is generally lower. The approaches belonging to this family compare fingerprints in term of features extracted from the ridge pattern. Two main reasons induced researchers to look for other fingerprint discriminative features beyond minutiae:

- Reliably extracting minutiae from poor quality fingerprints is very difficult. Although minutiae may carry most of the fingerprint discriminatory information, they do not always constitute the best trade-off between accuracy and robustness.
- Additional features may be used in conjunction with minutiae (and not as an alternative) to increase system accuracy and robustness.

The more commonly used alternative features are the size and shape of the fingerprint, number, type and position of singularities, spatial and geometrical attributes of the ridge lines, shape features, sweat pores, or global and local texture information.

3.1.2 System 1: The MorphoKit system

The purpose of the description of this commercial of-the-shelf system is to include a reference to a commercial product. As fingerprint technology is very mature, it is indeed important for the BEAT platform to be compatible with the use of commercial products. However there is no commitment to test the platform with the MorphoKit system during BEAT. This possibility will be investigated further on.

The MorphoKit is a software development kit (SDK) developed by Morpho which includes Morpho proprietary algorithms for generating minutiae templates and 1:1 or 1:N matching. A detailed description of the underlying algorithm cannot be disclosed and the following represents a brief specification of the SDK.

MorphoKit is a fingerprint acquisition and processing SDK. It is primarily intended to be used in the development of biometric application by private companies outside SagemDS, and will not include any of the specific features required for AFIS development (classification, segmentation of slap images, flatbed scanner management...) It is designed to be used in the development of small to medium scale biometric applications: it will included the best coding and 1:1 matching technology available today, but a limited version of Sagems 1:many matching technology, limited to small databases (3000 records for the standard product, no more than 100,000 records in any case), without the matching speed improvements specific to AFIS products. It is not designed for AFIS enrolment/identification or forensic applications and will not support 1000 dpi images or multi-finger images. The main features available in this SDK are the following :

- Coding of single-finger 500dpi grey-scale fingerprint images to create minutiae templates

- Authentication: 1:1 matching of single-finger fingerprint templates
- Identification: 1:many matching of reference template against a memory template database
- Template conversion of the proprietary CLV format to standard formats such as ANSI or ISO
- Live image acquisition with Morpho MSOXXX sensors

The fingerprint template format is CFV, which is a self-describing proprietary binary format. From the user's point of view, it is just a binary buffer of variable length (1300 bytes on average). Authentication and identification functions will only accept templates in CFV format. Matching templates is providing a score that can be compared to a given threshold for the decision: MATCH/NO MATCH. Reference thresholds are provided to meet a specific performance target. In other words, a target false accept rate (FAR) can be reached and guaranteed with a given fixed threshold.

3.1.3 System 2: The NFIS2 system

The minutiae-based NIST Fingerprint Image Software 2 (NFIS2) [26] is a minutiae-based fingerprint processing and recognition system formed from independent software, which constitutes a de facto standard reference system used in many fingerprint-related research contributions.

NFIS2 contains software technology, developed for the Federal Bureau of Investigation (FBI), designed to facilitate and support the automated manipulation and processing of fingerprint images. Source code for over 50 different utilities or packages and an extensive User's Guide are distributed on CD-ROM which is available free of charge[1].

From the 50 different software modules that are comprised within NFIS2, the most relevant for evaluation purposes are: MINDTCT for minutiae extraction, and BOZORTH3 for fingerprint matching.

MINDTCT

The MINDTCT system takes a fingerprint image and locates all minutiae in the image, assigning to each minutia point its location, orientation, type, and quality. The architecture of MINDTCT can be divided in the following stages: 1) Generation of image quality map; 2) Binarization; 3) Minutiae detection; 4) Removal of false minutiae, including islands, lakes, holes, minutiae in regions of poor image quality, side minutiae, hooks, overlaps, minutiae that are too wide, and minutiae that are too narrow (pores); 5) Counting of ridges between a minutia point and its nearest neighbours; 6) Minutiae quality assessment.

Because of the variation of image quality within a fingerprint, MINDTCT analyses the image and determines areas that are degraded. Several characteristics are measured, including regions of low contrast, incoherent ridge flow, and high curvature. These three

[1] `http://fingerprint.nist.gov/NFIS/`

conditions represent unstable areas in the image where minutiae detection is unreliable, and together they are used to represent levels of quality in the image. The image quality map of stage 1 is generated integrating these three characteristics. Images are divided into non-overlapping blocks, where one out of five levels of quality is assigned to each block.

The minutiae detection step scans the binary image of the fingerprint, identifying local pixel patterns that indicate the ending or splitting of a ridge. A set of minutia patterns is used to detect candidate minutia points. Subsequently, false minutiae are removed and the remaining candidates are considered as the true minutiae of the image. Fingerprint minutiae matchers often use other information in addition to just the points themselves. Apart from minutia's position, direction, and type, MINDTCT computes ridge counts between a minutia point and each of its nearest neighbours. In the last stage, MINDTCT assigns a quality/reliability measure to each detected minutia point. Even after performing the removal stage, false minutiae potentially remain in the list. Two factors are combined to produce a quality measure for each detected minutia point. The first factor is taken directly from the location of the minutia point within the quality map generated in stage 1. The second factor is based on simple pixel intensity statistics (mean and standard deviation) within the immediate neighbourhood of the minutia point. A high quality region within a fingerprint image is expected to have significant contrast that will cover the full grey-scale spectrum.

BOZORTH3

The BOZORTH3 matching algorithm computes a match score between the minutiae from any two fingerprints to help determine if they are from the same finger. It uses only the location and orientation of the minutiae points to match the fingerprints, and it is rotation and translation invariant. For fingerprint matching, compatibility between minutiae pairs of the two images are assessed by comparing the following measures: i) distance between the two minutiae and ii) angle between each minutia's orientation and the intervening line between both minutiae.

4 Iris biometrics

With the fast development of the iris image acquisition technology, iris recognition is expected to play a strong role in the future of biometric technology, with wide application areas in national ID cards, banking, e-commerce, welfare distribution, biometric passports and forensics, etc. Since the 1990s iris image processing and analysis research has achieved great progress. However, performance of iris recognition systems in unconstrained environments is still far from perfect. Iris localisation, nonlinear normalisation, occlusion segmentation, liveness detection, large-scale identification and many other research issues all need further investigation.

4.1 Systems

Iris recognition has become a popular research topic in recent years. Due to its reliability and nearly perfect recognition rates, iris recognition is used in high security areas. A literature review of the most prominent developed algorithms is showed below.

4.1.1 Existing systems

Looking at different approaches to analysing the texture of the iris has perhaps been the most popular area of research in iris biometrics. One of the first and most effective research lines has been the use of Gabor filters to produce a binary representation similar to Daugman's iriscode. Many different filters have been suggested for feature extraction. Sun et al. [28] use a Gaussian filter. Here the gradient vector field of an iris image is convolved with a Gaussian filter, yielding a local orientation at each pixel in the unwrapped template. Then the angle is quantised into six bins. This method was tested using an internal CASIA[2] dataset of 2,255 images obtaining an overall recognition rate of 100%. Another interesting approach with very good results is given by Monro et al. [29] where the discrete cosine transform is used for feature extraction. They apply the DCT to overlapping rectangular image patches rotated 45 degrees from the radial axis. The differences between the DCT coefficients of adjacent patch vectors are then calculated and a binary code is generated from their zero crossings. In order to increase the speed of the matching, the three most discriminating binarised DCT coefficients are kept, and the remaining coefficients are discarded.

Another research line focuses on using different types of filters to represent the iris texture with a real-valued feature vector. Ma et al. [30] use a variant of the Gabor filter at two scales to analyse the iris texture. They use Fisher's linear discriminant to reduce the original 1,536 features from the Gabor filters to a feature vector of size 200. Their experimental results show that the proposed method performs nearly as well as their implementation of Daugman's algorithm, and is a statistically significant improvement over other compared algorithms. The experimental results showed a correct recognition rate of 94.33% across the 2245 images of CASIA database.

[2] `http://www.cbsr.ia.ac.cn/IrisDatabase.htm`

It is also worth mentioning a third research line focused on finding effective ways to combine the previous two approaches. In this context, it should be highlighted the work of Hollingsworth et al. [31] where they acquire multiple iris codes from the same eye and evaluate which bits are the most consistent bits in the iriscode. They suggest masking the inconsistent bits in the iriscode to improve performance reaching equal error rates (EERs) of 0.068% under different subsets selected from the Iris Challenge Evaluation (ICE) dataset [27] using still images and video recordings.

4.1.2 System 1: VeriEye

VeriEye is a software development kit (SDK) developed by Neurotechnology which includes proprietary algorithms for generating minutiae templates and 1:1 or 1:N matching. A detailed description of the underlying algorithm cannot be disclosed and the following represents a brief specification of the SDK.

VeriEye is available as a software development kit that allows development of PC- and Web-based solutions on multiple platforms (Windows, Linux, Mac OS). It was released in 2008 by Neurotechnology and has been recognized by NIST as one of the most reliably accurate iris recognition algorithms in the 2009 Iris Exchange (IREX) independent evaluation campaign [35].

The proprietary algorithm implements advanced iris segmentation, enrollment and matching using robust digital image processing algorithms:

- Iris detection. The SDK includes algorithms to detect irises under non-ideal conditions such as images with obstructions, visual noise and different levels of illumination. The system also detects lighting reflections, eyelids and eyelashes obstructions. Images with narrowed eyelids or eyes that are gazing away are also accepted.

- Automatic interlacing detection and correction. Correction algorithms are applied to blurred images due to motion in order to improve the consistency of iris templates.

- Iris segmentation. VeriEye uses active shape models that more precisely model the contours of the eye, as iris boundaries are in general not modeled by perfect circles (or even ellipses). The centers of the iris inner and outer boundaries are different.

- Gazing-away eyes. Special segmentation algorithms are applied to iris images where the eye is detected to be in a non-frontal position with respect to the acquisition device.

- Matching. Configurable matching speed varies from 60,000 to 600,000 comparisons per second.

- Reliability. VeriEye 2.5 algorithm shows excellent performance when tested on all publicly available datasets. Especially good results are achieved on the NIST ICE2005 Exp1 database with iris images of intentionally degraded quality [27], and in the NIST IREX 2009 evaluation campaign [35].

4.1.3 System 2: Livor-Masek's system

The second iris recognition system to be used as initial reference system in the project was developed by L.M̃asek [32, 33]. The system takes two iris images as inputs, and gives a similarity score computed as the Hamming distance between the two binary templates of the input images (obtained applying Gabor filtering).

The system consists of the following sequence of steps: segmentation, normalisation, encoding and matching.

Segmentation
For the iris segmentation task, the system uses a circular Hough transform in order to detect the iris and pupil boundaries. Iris boundaries are modelled as two concentric circles. The range values where the radius should be searched is set manually. A maximum value is also imposed to the distance between the circles' centers. An eyelids and eyelashes removal step is also performed at this stage. Eyelids are isolated first by fitting a line to the upper and lower eyelid using the linear Hough transform. Eyelash isolation is then performed by histogram thresholding.

Normalisation
For normalisation of iris regions, a technique based on Daugman's rubber sheet model is employed. The center of the pupil is considered as the reference point, and radial vectors pass through the iris region. Since the pupil can be non-concentric to the iris, a remapping formula to rescale the center points depending on the angle around the circle is used. The normalisation stage produces: i) a 2D array with the angular resolution in the horizontal axis and the radial resolution in the vertical axis; ii) a 2D noise mask array for marking reflections, eyelashes, and eyelids detected in the segmentation stage.

Feature Encoding
Feature encoding is implemented by convolving the normalised iris pattern with 1D Log-Gabor wavelets. The 2D normalised pattern is broken up into a number of 1D signals, and then these ID signals are convolved with 1D Gabor wavelets. The rows of the 2D normalised pattern are taken as the 1D signals, each row corresponding to a circular ring on the iris region, where maximum independence occurs in [33]. The filtering output is then phase quantised to four levels using the Daugman method [34], with each filter producing two bits of data. The output of phase quantisation is a grey code, so that only 1 bit changes between adjacent quadrants. This will minimise the number of inconsistent bits in the case of two slightly misaligned intra-class patterns, and thus will provide a more accurate recognition rate [33]. The encoding process produces a bit-wise template containing a certain number of information bits, and a corresponding noise mask which represents corrupt areas within the iris pattern.

Matching
The Hamming distance (HD) is chosen as recognition metric, since bit-wise comparisons are necessary. The Hamming distance employed incorporates the noise mask, so that only significant bits are used for computing the matching score between two iris templates.

In order to account for rotational inconsistencies, when the Hamming distance of two templates is computed, one template is shifted left and right bit-wise and a number of Hamming distance values is obtained from each of the successive shifts [34]. This method corrects small misalignments in the normalised iris pattern caused by rotational differences during imaging. The lowest of the computed distance values is taken as the definitive similarity score.

5 Vein biometrics

Vascular Pattern or Vein Recognition is a relatively recent activity in biometrics, generally applied on hand, palm or finger, veins. It has only been recently that sensors to observe the vascular pattern in a convenient non-invasive way [36] have been commercialized. Similarly to fingerprints, the formation of the vascular network is governed by many different phenomena and it is widely accepted within the medical community that the vascular pattern is unique to each individual and quite stable with time.

For the BEAT project, we will consider the case of combined fingerprint and fingervein biometrics resulting from the use of *MorphoSmart^{TM}FINGER VP*.

5.1 Existing systems

Fingerprint systems are widely deployed and vein systems have now several operational applications. As described in Section 3, there are various techniques for fingerprint recognition. Similarly, researchers are exploring different solutions for vein recognition. The work by Hartung et al. [38] is a good example of a specific recognition algorithm. In this work we may also find a complete survey on the different feature extraction algorithms that have been proposed in the literature.

However the use of a single device to combine fingerprint and fingervein restricts us to very few existing solutions. Among these there are a NEC fingerprint and fingervein sensor and the MorphoSmart^{TM}FINGER VP device. And Fujitsu recently announced the development of a future fingerprint and palm-vein sensor.

5.1.1 System 1: MorphoSmart^{TM}FINGER VP

The Morpho FINGER VP system has been developed for the MorphoSmart^{TM}FINGER VP device. This solution combines Hitachi's vein imaging technology to detect the pattern of blood vessels under the skin and Morpho's fingerprint identification technology.

From a pure vein biometrics perspective, the technology implemented in the device is based upon patented technology developed by Hitachi-Omron, the fundamentals of which are described in [37].

The basic principle is to select an illumination wavelength for which absorption from deoxidized hemoglobin (flowing freely in the blood stream) will be maximized while the "background" absorption (all other cell tissues) will be minimal. This way the vascular pattern will appear in great contrast "through" the different skin layers in the finger. It is also interesting to note that only veins – and not arteries – are going to be visible: Arteries carry much less deoxidized hemoglobin than veins; Arteries lie deeper in the finger than veins.

The sensor by illuminating a finger with near infra-red wavelengths light allows to build a reconstitution of the network depending on the recorded absorption. A camera collects the network projection. Once the vascular network is captured, the acquired image is then processed through standard image processing techniques to enhance the relevant signal

and diminish noise, down to a smaller number of gray levels to be able to perform efficient matching.

Thanks to the extraction of specific features, the vein pattern together with the captured fingerprint are then compared to all templates stored in the database. The result of this search query can be either positive (the person searched for was indeed present in the database), or negative (the person was not found in the database). A second mode of operation is also possible: given a stored template of a finger, the system can "verify" it against a live finger. This can be used as a verification of a claimed identity.

Furthermore, the Morpho FINGER VP system performs an enhanced fusion of fingerprint and fingervein comparisons. This dedicated fusion allows to maximize accuracy and guarantees the chosen security level (FAR).

The system is designed to work on sets of images (fingervein and fingerprint) acquired from FINGER VP device. From these images, a template is created in a proprietary format. The template matching algorithm provides a consolidated score that can be compared to a given threshold in order to obtain a decision. Reference thresholds are provided to meet a specific performance target, i.e. a target False Accept Rate can be reached and guaranteed with a given fixed threshold. The main features available in this system are the following:

- Encoding of single-finger 500dpi grey-scale fingerprint images and fingervein images to create templates of selected features;
- Authentication (or verification): 1-to-1 comparison of fingerprint + fingervein templates;
- Identification: 1-to-many comparison of reference template against a database.

The system only outputs the final result and there is no access to independent decision results for each modality.

6 Multimodal biometrics

6.1 Classification of fusion algorithms for multi-system biometrics

Fusion algorithms can be categorised into pre-classification and post-classification strategies. By pre-classification, we understand that the underlying information sources are combined prior to invoking a classifier. An example is combining several images together to form a larger image. This is known as *sensor-level fusion*. Biometric data can also be combined at the feature level, for instance, combining two types of representation, that is, shape and texture in 3D face recognition. This is known as *feature-level fusion*.

By far, most fusion algorithms operate after the classification stage, hence, belong to the *post-classification* strategy. Information at this level are hypotheses about the identity of the subjects, which can be matching scores, a list of probable identities or rank, and hard decisions derived from the matching scores. The different fusion algorithms, along with their representative methods, are shown in Figure 1.

For the rest of the discussion, we shall categorise fusion algorithms by their context of usage, that is, either for authentication (verification) or for closed-set/open-set identification. This approach is more appropriate for the BEAT project because performance metrics and fusion algorithms are very different for both application types.

6.2 Fusion algorithms for authentication

The task of a fusion module in biometric authentication is to determine whether an identity claim is genuine or not. This is essentially a binary classification task. This implies that one can use any binary classifier for this purpose, for instance, logistic regression, neural network and support vector machines. This is known as the *classification approach*.

An althernative strategy takes advantage of the fact that the output of biometric matcher (matching score) is associated with a hypothesis. For instance, high similarity scores imply a genuine idedntity claim whereas low similarity scores imply an impostor claim. Distance metrics is interpretted exactly in the opposite way. The nature of matching scores being assertive about an identity claim is exploited by an alternative approach to fusion, namely, the *combination approach*. In this section, we will present the combination approach followed by the classification approach. The final section gives more insight into other more advanced fusion methods reported in the literature.

6.2.1 Combination approach

In the combiantion approach, one has to calibrate the matching scores because each biometric matcher can produce outputs in different range and their respective score distribtuions, either produced by genuine users or impostors, are very different. Consequently, the combination approach requires two steps in the fusion process, namely, a score normalisation step followed by a combination rule.

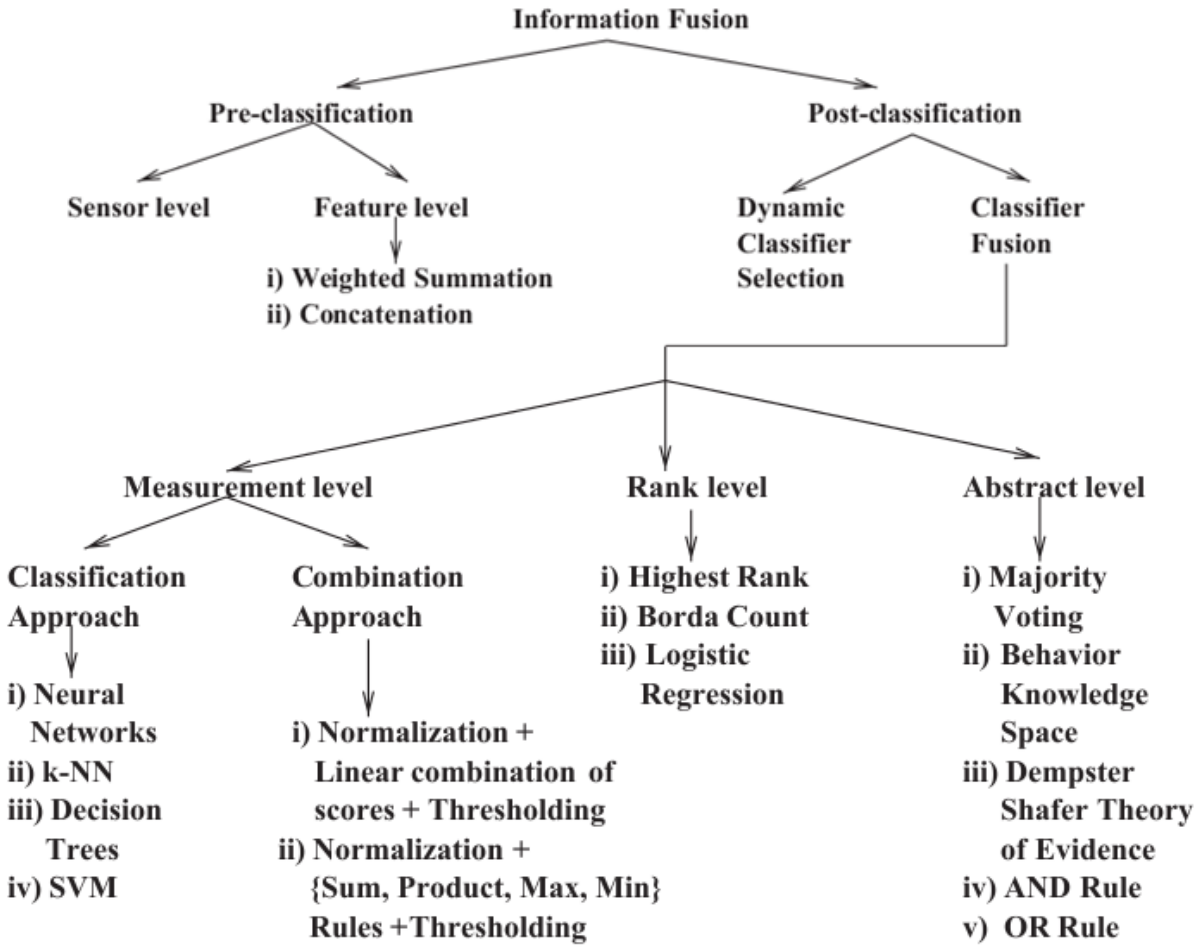

Figure 1: Categorisation of algorithms for biometric information fusion

Jain *et al* [39] investigated and compared different score normalisation techniques. They are:

- **min-max normalisation**:

$$f_{mm}(y) = \frac{y - \min_{y' \in \mathcal{Y}}}{\max_{y' \in \mathcal{Y}} - \min_{y' \in \mathcal{Y}}}$$

- **median absolute deviation normalisation**:

$$f_{mad}(y) = \frac{y - \text{median}_{y' \in \mathcal{Y}}}{MAD}$$

where

$$MAD = \text{median}_{y \in \mathcal{Y}}(|y - \text{median}_{y' \in \mathcal{Y}}|)$$

- **Z-score normalisation**:

$$f_z(y) = \frac{y - \mu}{\sigma}$$

- **normalisation via logistic regression**:

$$f_{LR}(y) = yw_1 + w_o$$

where the weight parameters are obtained after training a logistic regression of the form:

$$P(\omega_1|y) = \frac{1}{1 + \exp(-f_{LR}(y))}$$

In the above notation, $y \in \mathcal{Y}$ implies that the matching score y is taken from the training set $\mathcal{Y}$. μ and σ are the mean and standard deviation of the system score derived from the training set $\mathcal{Y}$.

Having normalised the scores so that the scores of different biometric matchers are in comparable range, the next step is to combine them using geometric operators such as min, max, sum, and product rules.

Let us denote the vector of biometric matchers to be $\mathbf{y}$ and its element y_m for $m \in \{1, \ldots, M\}$ and m is the indice to one of the biometric matchers. We will also use f_{norm} to denote any one of the score normalisation schemes described above, i.e., $norm \in \{mm, mad, z\}$. The fixed fusion rules can then be written as:

$$\begin{aligned}
f_{min}(\mathbf{y}) &= \min_{m=1}^{M} f_{norm}(y_m) \\
f_{max}(\mathbf{y}) &= \max_{m=1}^{M} f_{norm}(y_m) \\
f_{sum}(\mathbf{y}) &= \sum_{m=1}^{M} f_{norm}(y_m) \\
f_{prod}(\mathbf{y}) &= \prod_{m=1}^{M} f_{norm}(y_m)
\end{aligned}$$

Among these rules, the sum rule often give reasonably good performance [44]. This is because the noise variance from individual biometric matchers are reduced due to the sum operator. If there are M biometric matchers, it can be shown that the noise variance of the compound (fused) score, conditioned on each class (genuine or impostor score distributions), can be reduced by at most $\frac{1}{M}$. Therefore, fusion almost always reduces the classification error when it is carried out correctly. Although one can show that the performance of a fusion system is *always* better than the "average" performance of the individual matcher, it is not guaranteed that the fusion performance is always better than the best biometric matcher in the pool. By assuming that the class-conditional scores to be normally distributed, Poh and Bengio [45] proposed a predictive fusion model in terms of Equal Error Rate that studies the condition required for a fusion system to be as good as the best biometric matcher in the pool.

6.2.2 Classification approach

In theory, many off-the-shelf classifiers such as neural network, support vector machines, and decision trees can be used to combine scores from different biometric matchers. One can categorise these methods by the output type (posterior probability, log-likelihood ratio, or margin), optimisation method and the criterion choice (enforcing sparsity, regularisation, brute force, genetic algorithms), and their nature (being generative or discriminative). For the purpose of our discussion, we shall use the generative/discriminative dichotomy and the output types, as well as include a class of weighted sum heuristic methods. The key inclusion criteria for the fusion algorithms discussed here are their relevance and evidence of their effectiveness based on the literature.

Weighted sum heuristics

Weighted sum fusion has a generic form:

$$f_{wsum}(\mathbf{y}) = \sum_{m=1}^{M} f_{norm}(y_m) w_m$$

where f_{norm} is any score normalisation scheme. In theory, the score normalisation is optional but in practice, by mapping scores into a common range, it is easier to find the weight parameters.

There are a number of schemes that attempt to find weights. The general idea is that the more accurate biometric matcher should be given a higher weight. Based on this intuition, the weight should be *inversely proportional* to the error or uncertainty, or be proportional to a separability criterion. The following are some examples:

- Equal weight:
$$w_m = \frac{1}{M} \text{ for all } m$$
- EER as a function of weight:
$$w_m = \frac{(\text{EER}_m)^{-1}}{\sum_{m'=1}^{M} (\text{EER}_{m'})^{-1}}$$
- d-prime as a function of weight:
$$w_m = \frac{d'_m}{\sum_{m'=1}^{M} d'_m}$$
where
$$d' = \frac{\mu_1 - \mu_0}{\sqrt{\sigma_1^2 + \sigma_2^2}}$$

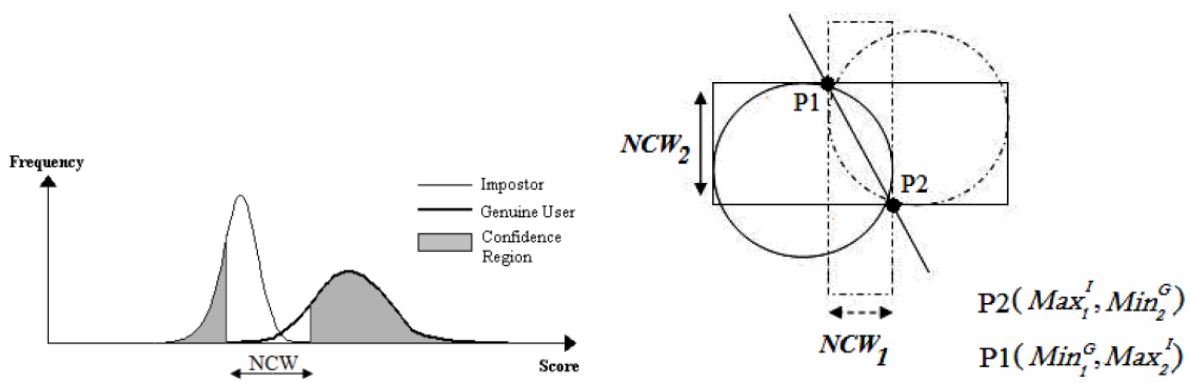

(a) NCW (b) Rationale for NCW

Figure 2: Non-confidence width

- Non-confidence width as a function of weight:

$$w_m = \frac{(\text{NCW}_m)^{-1}}{\sum_{m'=1}^{M}(\text{NCW}_{m'})^{-1}}$$

where

$$\text{NCW}_m = \max_{y \in \mathcal{Y}_{m,0}}(y) - \min_{y \in \mathcal{Y}_{m,0}}(y)$$

In the above equations, μ_ω is the mean and ω is the standard deviation of one of the two possible matching events, that is, genuine when $\omega = 1$ and impostor when $\omega = 0$.

In the NCW equation, $\mathcal{Y}_{m,0}$ denotes the matching scores conditioning on being genuine or impostor, that is, $\mathcal{Y}_{m,\omega} \equiv \{y | y \in \{y | y \in \mathcal{Y}_m, \omega\}$. Effectively, NCW defines the length of the region in which the training scores make either fales acceptance or false rejection error, as depicted in Figure 2(a). The rationale for NCW is that two points $P1$ and $P2$ that define the intersection of two circles (representing genuine and impostor scores) is an ideal separable decision plan between them. This idea is shown in Figure 2(b) that is applicable to the case of combining two biometric matchers. The dashed circle represents the distribution of genuine scores whereas the circle plotted with a continuous line represents that of impostor scores. Although NCW does not take into account the correlation among the biometric matcher outputs, Chia *et al*[40] reported that its performance is better than the remaining heuristics reported here.

Generative classifiers

Bayesian classifiers have a generic form that exploits the Bayes rule. It relies on the estimate of the class-conditional densities $p(\mathbf{y}|\omega)$ and the prior probability $P(\omega)$ for the impostor or genuine matching scores, $\omega \in \{0, 1\}$, respectively. There are two types of output that can be produced by a Bayesian classifier, namely in terms of posterior probability,

$$f_{prob}(\mathbf{y}) = P(\omega_1|\mathbf{x}) = \frac{p(\mathbf{y}|\omega_1)P(\omega_1)}{\sum_\omega p(\mathbf{y}|\omega)P(\omega)}$$

or log-likelihood ratio in accordance to the Neyman-Pearson lemma:

$$f_{llr}(\mathbf{y}) = \log \frac{p(\mathbf{y}|\omega = 1)}{p(\mathbf{y}|\omega = 0)}.$$

In the case of posterior probability, the decision threshold is 0.5 but the prior probabilities $P(\omega)$ has to be set accordingly. In the case of log-likelihood ratio, the classifier output is compared with a decision threshold which is an implicit function of the prior probabilities. Both methods will give exactly the same result if the underlying class-conditional densities are the same.

The main interest in this Bayesian classifier is its flexibility: one can use any density function in order to estimate $p(\mathbf{y}|\omega)$. Examples are

- Gaussian Mixture Models (GMM), leading to a GMM-Bayes classifier,
- Gaussian models, leading to a Gaussian classifier, and
- a product of marginal densities

$$p(\mathbf{y}|\omega) = \prod_m (y_m|\omega)$$

leading to a Naive Bayes classifier.

Other density estimators that have been used are kernel density estimator or Parzen window [46] and Gaussian copula [47].

Discriminative classifier estimating posterior probability

Logistic regression is arguably the most commonly used discriminative classifier owing to its simplicity: a linear classifier with convex optimisation, hence, leading to a global solution. Although non-linear classifiers have been used, they often do not offer significantly better performance [41].

Logistic regression is defined as:

$$y_{com} \equiv P(\mathtt{C}|\mathbf{x}) = \frac{1}{1 + \exp(-g(\mathbf{x}))}, \tag{1}$$

where

$$g(\mathbf{x}) = \sum_{j=1}^{M} \beta_j x_j + \beta_0, \tag{2}$$

where x_j's are elements in $\mathbf{x}$. The weight parameters β_j are optimized using gradient ascent to maximize the likelihood of the training data given the logistic regression model [43]. It can be shown that the following relationship is true:

$$g(\mathbf{x}) = \log \frac{P(\mathtt{C}|\mathbf{x})}{P(\mathtt{I}|\mathbf{x})}. \tag{3}$$

6.2.3 Other fusion algorithms and settings

The brief discussion is certainly not exhaustive. For instance, apart from classifiers constructed through the Bayes rule, one can also use the Dempster-Shafer Theory [48].

Other fusion settings include

- **Client-specific fusion**, where a fusion classifier is tailored to each cliamed identity [49].
- **Quality-based fusion**, where a fusion classifier weighs the output of each biometric matcher as a function of the signal quality [50].
- **Sequential fusion**, where the matching scores to be combined are considered sequentially, until a certain level of confidence is reached, or all the match scores are exhausted [51].

6.3 Fusion algorithms for identification

We shall also conveniently describe fusion algorithms for identification in two categories: fixed rule rank-level fusion and its trainable counterpart.

6.3.1 Fixed rule rank-level fusion

In identfication, the most commonly used information for fusion is the rank of the identity list. Let $r_{i,m}$ denote rank i of matcher m and there are $i = 1, \ldots, N$ identities in the database. Several rank-level fusion methods have been proposed in the literature [52]; they are listed below:

- **Highest rank method**:
$$R_j = \min_{i=1}^{N} r_{i,m}$$
- **Borda count method**:
$$R_j = \sum_{i=1}^{N} r_{i,m}$$

The highest rank method will often result in a tie, a situation where several fused ranks have the same value. Ties are broken randomly.

It is interesting to note that the highest rank method is a min rule fusion whereas the Borda count method is a sum rule fusion. Although these operators are the same as those applied to the combination approach, the domain of application is different. Here, the rank information is used whereas in the combination rule for authentication, the score or measurement is used.

6.3.2 Trainable rank-level fusion

It is possible to extend these rules in order to add some degree of flexibility. For instance, one can weight the rank prior to computing the min:

$$R_j = \min_{i=1}^{N}(w_i \times r_{i,m})$$

where w_i is the weight associated with each rank.

One can also add non-linearity to the output in a way that resembles the perceptron algorithms:

$$R_j = \frac{1}{1 + \exp\left(-\sum_{i=1}^{N} w_i r_{i,m}\right)},$$

$$R_j = \sum_{i=1}^{N} \tanh(w_i \times r_{i,m}),$$

and

$$R_j = \sum_{i=1}^{N} w_i \exp(r_{i,m})$$

The first formula is a logistic regression, the weight of which can be trained using "gradient ascent". The remaining two methods can be optimised using general purpose optimisation such as the ant colony optimisation algorithm and its variants [53].

In all of the methods discussed in this section, the weight parameters w_i have to be tuned based on a development set. This means that this data set should not be used for testing the performance of the algorithm; a sepeate held-out test set should be used for this purpose. A second issue regarding trainable rank-level fusion is how to calculate the weight when an additional subject is added to the gallary.

Trainable rank-level fusion is postulated to outperform its simple fixed rule fusion only when the performance of the underlying biometric matchers are very different from each other. This is because the weights can be used to adjust the contribution of each biometric matcher to the final combined rank list.

6.4 Initial reference fusion algorithms

To summarise, in this section, we have surveyed a number of fusion algorithms and settings for biometric authentication and identification. For each application type, fixed rule fusion and trainable fusion algorithms are discussed. As initial reference fusion systems, however, we will opt for algorithms that have been shown to be robust, simple, yet perform well. For the authentication scenario, we will opt for the sum rule with scores calibrated using logistic regression or min-max score score normalisation. For the identification scenario, we will consider the highest rank method which has been reported to be very effective.

7 Summary

This document describes the different biometric systems that will be considered within BEAT for performance, security and countermeasure evaluation purposes.

The systems have been chosen in order to provide for reliable, meaningful results which are easily compared with other previous works with minimal effort being spent on adapting existing systems to new data and protocols. This will allow us to concentrate effort on the definition of new objective and replicable metrics, protocols and benchmarks for the evaluation of performance, security and countermeasures which form the heart of the BEAT project. In some cases, however, particularly for the more experimental biometrics, not all the systems described in the present document will be necessarily used in all the evaluation tests. In general, the systems used for a particular experimental setup will be fixed on a case by case basis.

The biometric systems considered here are those used by each partner in previous work and are state-of-the-art. The document outlines each biometric system with a brief account of the setup in each case. No evaluation work is reported in this document.

Also reported are the fusion systems that will be used for multi-modal work, together with an account of the state-of-the-art in data fusion.

References

[1] J. Galbally, C. McCool, J. Fierrez, S. Marcel and J. Ortega-Garcia, On the Vulnerability of Face Verification Systems to Hill-Climbing Attacks, Pattern Recognition, Volume 43, pp. 1027-1038, 2010.

[2] R. Jafri, Hamid R. Arabnia, A Survey of Face Recognition Techniques, Journal of Information Processing Systems, Volume 5, No. 2, pp. 41-64, June, 2009.

[3] F. Cardinaux, C. Sanderson, and S. Marcel. Comparison of MLP and GMM classifiers for face verification on XM2VTS. In Proc. International Conference on Audio- and Video-based Biometric Person Authentication, pages 10581059, 2003.

[4] S. Lucey and T. Chen. A GMM parts based face representation for improved verification through relevance adaptation. In Proc. International Conference on Computer Vision and Pattern Recognition (CVPR), volume 2, pages 855861, 2004.

[5] D. A. Reynolds. Comparison of background normalization methods for text independent speaker verification. In Proc. European Conference on Speech Communication and Technology (EuroSpeech), pages 963966, Rhodes, Greece, September 1997.

[6] [215] W. Zhao, R. Chellappa, P. Phillips, and A. Rosenfeld. Face recognition: A literature survey. ACM Computing Surveys, 35(4):399459, 2003.

[7] K. R. Rao and P. Yip, Discrete Cosine Transform: Algorithms, Advantages, Applications, Academic Press, 1990.

[8] A. Anjos et al., Bob: a free signal processing and machine learning toolbox for researchers, submitted to the ACM Multimedia Conference, Japan, 2012.

[9] S. Prince, J. H. Elder, J. Warrell and F. Felisberti, Tied factor analysis for face recognition across large pose differences, IEEE Pattern Analysis and Machine Intelligence, 2008.

[10] N. Kumar, A. Berg, N. Belhumeur and S. Nayar, Attribute and Simile Classifiers for Face Verification, Proceedings of the 2009 International Conference on Computer Vision, 2009.

[11] C. Sanderson and K. K. Paliwal, Fast feature extraction method for robust face verification, Electronic Letters, Number 25, Volume 38, pp. 1648-1650, 2002.

[12] G. Heusch, Y. Rodriguez and S. Marcel, Local Binary Patterns as an Image Preprocessing for Face Authentication, International Conference on Automatic Face and Gesture Recognition, pp. 9-14, 2006.

[13] F. Cardinaux, C. Sanderson and S. Bengio, User Authentication via Adapted Statistical Models of Face Images, IEEE Trans. Signal Processing, pp. 681-685, 2006.

[14] Timo Ojala, Matti Pietikäinen and Topi Mäenpää, Multiresolution Gray-Scale and Rotation Invariant Texture Classification with Local Binary Patterns, IEEE Trans. Pattern Anal. Mach. Intell., Volume 24, Number 7, pp. 971-987, 2002.

[15] C. Chan and J. Kittler and K. Messer, Multi-scale Local Binary Pattern Histograms for Face Recognition, Proceedings of the 2007 International Conference on Biometrics, pp. 809-818, 2007.

[16] W. Zhao and R. Chellappa and P. Phillips and A. Rosenfeld, Face Recognition: A Literature Survey, ACM Computing Surveys, pp. 399-459, Volume 35, Number 4, 2003.

[17] Y. Rodriguez and S. Marcel, Face Authentication Using Adapted Local Binary Pattern Histograms, Proceedings of the 2006 European Conference on Computer Vision, pp. 321-332, 2006.

[18] W. Zhang, S. Shan, W. Gao, X. Chen and H. Zhang, Local Gabor Binary Pattern Histogram Sequence (LGBPHS): A Novel Non-Statistical Model for Face Representation and Recognition, Proceedings of the 2005 International Conference on Computer Vision, pp. 786-791, 2005.

[19] T. Ahonen and M. Pietikäinen, Pixelwise Local Binary Pattern Models of Faces Using Kernel Density Estimation, Proceedings of the Third International Conference on Advances in Biometrics, pp. 52-61, 2009.

[20] T. Ahonen, A. Hadid and M. Pietikäinen, Face Recognition with Local Binary Patterns, Proceedings of the 2004 European Conference on Computer Vision, pp. 469-481, 2004.

[21] `http://idiap.github.com/bob`

[22] D. Maltoni, D. Maio, A.K. Jain and S. Prabhakar, Handbook of Fingerprint Recognition (Second Edition), Springer (London), 2009.

[23] B.G. Sherlock, D.M. Monro and K. Millard, Algorithm for enhancing fingerprint images, Electronics Letters, vol. 28, no. 18, pp. 1720, 1992.

[24] O. Trier and A.K. Jain, Goal-directed evaluation of binarization methods, IEEE Transactions on Pattern Analysis Machine Intelligence, vol. 17, no. 12, pp. 1191-1201, 1995.

[25] A. M. Bazen and S.H. Gerez, Fingerprint matching by thin-plate spline modelling of elastic deformations, Pattern Recognition, vol. 36, no. 8, pp. 1859-1867, 2003.

[26] M. Garris, C. Watson, R. McCabe, and C. Wilson. User's guide to NIST Fingerprint Image Software 2 (NFIS2). National Institute of Standards and Technology, 2004.

[27] National Insitute of Standards and Technology. Iris challenge evaluation data. http://iris.nist.gov/ice/.

[28] Z. Sun, T. Tan, and Y. Wang. Robust encoding of local ordinal measures: A general framework of iris recognition. In Proc. BioAW Workshop, pp. 270-282, 2004.

[29] D. Monro, S. Rakshit, and D. Zhang. DCT-based iris recognition. IEEE Transactions on Pattern Analysis and Machine Intelligence, Vol. 29, No.4, pp.586-595, April 2007.

[30] L. Ma, T. Tan, Y. Wang, and D. Zhang. Personal identification based on iris texture analysis. IEEE Transactions on Pattern Analysis and Machine Intelligence, Vol. 25, No.12, pp. 1519-1533, 2003.

[31] K. Hollingsworth, K. Bowyer, and P. Flynn. All iris code bits are not created equal. In Biometrics: Theory, Applications, and Systems, Sept 2007.

[32] L. Masek and P. Kovesi, Matlab source code for a biometric identification system based on iris patterns, The School of Computer Science and Software Engineering, The University of Western Australia, 2003.

[33] L. Masek et al., Recognition of human iris patterns for biometric identification, Ph.D. thesis, Technical Report, The school of Computer Science and Software Engineering, The University of Western Australia, 2003.

[34] J. Daugman, How iris recognition works, Proceedings of 2002 International Conference on Image Processing, vol. 1, pp. 22-25, 2002.

[35] P. Grother, E. Tabassi, G. W. Quinn, and W. Salamon. IREX I: Performance of iris recognition algorithms on standard images, National Institute of Standards and Technology, Tech. Rep., 2009.

[36] K. Shimizu, Optical trans-body imaging - Feasibility of optical CT and Functional Imaging of Living Body, Medicina Philisophica, 11:620-629, 1992.

[37] Mitsutoshi Himaga and Katsuhiro Kou, Finger Vein Authentication technology and financial applications, Advances in Biometrics, 1, 89-105, 2008.

[38] Daniel Hartung, Anika Pflug and Christoph Busch, Feature extraction from vein images using spatial information and chain codes, Information Security Technical Report, Volume 17, Issues 12, February 2012, Pages 2635, 2012.

[39] A. Jain, K. Nandakumar, A. Ross (2005). Score Normalisation in Multimodal Biometric Systems. Pattern Recognition 38(12): 2270-2285.

[40] Chia, C., Sherkat, N., and Nolle, L. (2010). Towards a best linear combination for multimodal biometric fusion. In Proceedings of the 2010 20th International Conference on Pattern Recognition, ICPR '10, pages 1176-1179, Washington, DC, USA. IEEE Computer Society.

[41] Non-Linear Variance Reduction Techniques in Biometric Authentication, Workshop on Multimodal User Authentication (MMUA), 2003.

[42] A.P. Dempster, "A generalization of Bayesian inference," J. Royal Statist. Soc., ser. B, vol. 30, pp. 205–247, 1968.

[43] A. J. Dobson, An Introduction to Generalized Linear Models, CRC Press, 1990.

[44] J. Kittler, M. Hatef, R. P.W. Duin, and J. Matas, "On Combining Classifiers," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 20, no. 3, pp. 226–239, 1998.

[45] N. Poh and S. Bengio, "How Do Correlation and Variance of Base Classifiers Affect Fusion in Biometric Authentication Tasks?," *IEEE Trans. Signal Processing*, vol. 53, no. 11, pp. 4384–4396, 2005.

[46] E. Parzen "On estimation of a probability density function and mode". Annals of Mathematical Statistics 33: 1065–1076, 1962.

[47] Dass, S. C., Nandakumar, K. and Jain, A. K., "A Principled Approach to Score Level Fusion in Multimodal Biometric Systems", Proc. of Audio- and Video-based Biometric Person Authentication (AVBPA) 2005, pp. 1049-1058, Rye Brook, NY, July 2005.

[48] N. Poh, T. Bourlai, J. Kittler, L. Allano, F. Alonso-Fernandez, O. Ambekar, J. Baker, B. Dorizzi, O. Fatukasi, J. Fierrez, H. Ganster, J. Ortega-Garcia, D. Maurer, A. A Salah, T. Scheidat, and C. Vielhauer, Benchmarking Quality-dependent and Cost-sensitive

Multimodal Biometric Fusion Algorithms, IEEE Trans. on Information Forensics and Security, 4(4), pp. 849–866, 2009.

[49] N. Poh and J. Kittler, Incorporating Variation of Model-specific Score Distribution in Speaker Verification Systems, IEEE Trans. Audio, Speech and Language Processing, 16(3):594-606, 2008.

[50] N. Poh and J. Kittler, A Unified Framework for Biometric Expert Fusion Incorporating Quality Measures, IEEE Trans. on Pattern Analysis and Machine Intelligence, 34(1):3-18, 2012.

[51] L. Allano, B. Dorizzi and S. Garcia-Salicetti, Tuning cost and performance in multi-biometric systems: A novel and consistent view of fusion strategies based on the Sequential Probability Ratio Test (SPRT), Pattern Recognition Letters, Vol. 31, No. 9, pp. 884-890, 2010.

[52] M. M. Monwar, M. L. Gavrilova, Multimodal biometric system using rank-level fusion approach. IEEE transactions on systems, man, and cybernetics. Part B, Cybernetics : a publication of the IEEE Systems, Man, and Cybernetics Society, Vol. 39, No. 4, pp. 867-878, 2009.

[53] K. Veeramachaneni, L. A. Osadciw, P. K. Varshney, Adaptive Multimodal Biometric Fusion Algorithm Using Particle Swarm in SPIE Aerosense, 2003