Chapter 11
Multibiometric Systems: Overview, Case Studies and Open Issues

Arun Ross and Norman Poh

Abstract Information fusion refers to the reconciliation of evidence presented by multiple sources of information in order to generate a decision. In the context of biometrics, evidence reconciliation plays a pivotal role in enhancing the recognition accuracy of human authentication systems and is referred to as multibiometrics. Multibiometric systems combine the information presented by multiple biometric sensors, algorithms, samples, units, or traits in order to establish the identity of an individual. Besides enhancing matching performance, these systems are expected to improve population coverage, deter spoofing, facilitate continuous monitoring and impart fault-tolerance to biometric applications. This chapter introduces the topic of multibiometrics and enumerates the various sources of biometric information that can be consolidated as well as the different levels of fusion that are possible in a biometric system. The role of using ancillary information such as biometric data quality and soft biometric traits (e.g., height) to enhance the performance of these systems is discussed. Three case studies demonstrating the benefits of a multibiometric system and the factors impacting its architecture are also presented. Finally, some of the open challenges in multibiometric system design and implementation are enumerated.

11.1 Introduction

A reliable identity management system is a critical component in several applications that render services to only legitimately enrolled users. Examples of such applications include sharing networked computer resources, granting access to nuclear

Arun Ross
West Virginia University, Morgantown, West Virginia, USA. e-mail: arun.ross@mail.wvu.edu

Norman Poh
University of Surrey, Guildford, GU2 7XH, Surrey, UK. e-mail: normanpoh@ieee.org

facilities, performing remote financial transactions or boarding a commercial flight. The proliferation of web-based services (e.g., online banking) and the deployment of decentralized customer service centers (e.g., credit cards) have further enhanced the need for reliable identity management systems. Traditional methods of establishing a person's identity include knowledge-based (e.g., passwords) and token-based (e.g., ID cards) mechanisms, but these surrogate representations of identity can be easily lost, shared, manipulated or stolen thereby undermining the intended security. Biometrics offers a natural and reliable solution to certain aspects of identity management by utilizing fully automated or semi-automated schemes to recognize individuals based on their inherent physical and/or behavioral characteristics [28]. By using biometrics (see Fig. 11.1) it is possible to establish an identity based on *who you are*, rather than by *what you possess*, such as an ID card, or *what you remember*, such as a password.

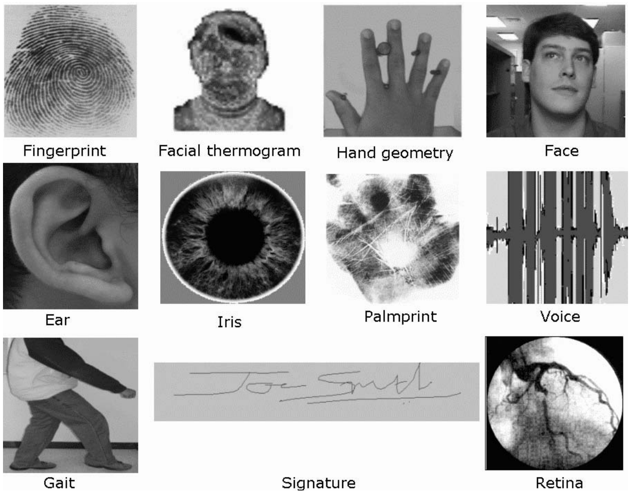

Fig. 11.1 Examples of biometric traits that can be used for authenticating an individual.

Most biometric systems that are presently in use, typically use a single biometric trait to establish identity (i.e., they are unibiometric systems). Some of the challenges commonly encountered by biometric systems are listed below:

1. Noise in sensed data: The biometric data being presented to the system may be contaminated by noise due to imperfect acquisition conditions or subtle variations in the biometric itself. For example, a scar can change a subject's fingerprint while the common cold can alter the voice characteristics of a speaker. Similarly, unfavorable illumination conditions may significantly affect the face and iris images acquired from an individual. Noisy data can result in an individual being incorrectly labeled as an impostor thereby increasing the False Reject Rate (FRR) of the system.

2. Non-universality: The biometric system may not be able to acquire meaningful biometric data from a subset of individuals resulting in a failure-to-enroll (FTE) error. For example, a fingerprint system may fail to image the friction ridge structure of some individuals due to the poor quality of their fingerprints. Similarly, an iris recognition system may be unable to obtain the iris information of a subject with long eyelashes, drooping eyelids or certain pathological conditions of the eye. Exception processing will be necessary in order to include such users into the authentication system.
3. Upper bound on identification accuracy: The matching performance of a unibiometric system cannot be continuously improved by tuning the feature extraction and matching modules. There is an implicit upper bound on the number of distinguishable patterns (i.e., the number of distinct biometric feature sets) that can be represented using a template. The capacity of a template is constrained by the variations observed in the feature set of each subject (i.e., *intra*-class variations) and the variations between feature sets of different subjects (i.e., *inter*-class variations).
4. Spoof attacks: Behavioral traits such as voice [15] and signature [21] are vulnerable to spoof attacks by an impostor attempting to mimic the traits corresponding to legitimately enrolled subjects. Physical traits such as fingerprints can also be spoofed by inscribing ridge-like structures on synthetic material such as gelatine and play-doh [41,55]. Targeted spoof attacks can undermine the security afforded by the biometric system and, consequently, mitigate its benefits [56].

Some of the limitations of a unibiometric system can be addressed by designing a system that consolidates *multiple* sources of biometric information. This can be accomplished by fusing, for example, multiple traits of an individual, or multiple feature extraction and matching algorithms operating on the same biometric. Such systems, known as multibiometric systems [27, 60], can improve the matching accuracy of a biometric system while increasing population coverage and deterring spoof attacks.

The rest of the chapter is structured as follows. Section 11.2 lists some of the advantages of using multibiometric systems; Section 11.3 presents the taxonomy used to characterize these systems; Section 11.4 provides an overview of the various levels of fusion that are possible; Section 11.5 discusses the possibility of incorporating ancillary features such as soft biometrics and quality in order to enhance the matching performance of fusion systems; Section 11.6 presents three case studies highlighting the benefits of fusion and the factors impacting its architecture; Section 11.7 lists some of the open challenges in biometric fusion; Section 11.8 summarizes the contributions of this chapter.

11.2 Advantages of Multibiometric Systems

Besides enhancing matching accuracy, the other advantages of multibiometric systems over traditional unibiometric systems are enumerated below [60].

1. Multibiometric systems address the issue of non-universality (i.e., limited population coverage) encountered by unibiometric systems. If a subject's dry finger prevents her from successfully enrolling into a fingerprint system, then the availability of another biometric trait, for instance her iris, can aid in the inclusion of the individual in the biometric system. A certain degree of flexibility is achieved when a user enrolls into the system using several different traits (e.g., face, voice, fingerprint, iris, hand) while only a subset of these traits (e.g., face and voice) is requested during authentication based on the nature of the application under consideration and the convenience of the user.
2. Multibiometric systems can facilitate the filtering or indexing of large-scale biometric databases. For example, in a bimodal system consisting of face and fingerprint, the face feature set may be used to compute an index value for extracting a candidate list of potential identities from a large database of subjects. The fingerprint modality can then determine the final identity from this limited candidate list.
3. It becomes increasingly difficult (if not impossible) for an impostor to spoof multiple biometric traits of a legitimately enrolled individual. If each sub-system indicates the probability that a particular trait is a 'spoof', then appropriate fusion schemes can be employed to determine if the user is, in fact, an impostor or not. Furthermore, by asking the user to present a random subset of traits at the point of acquisition, a multibiometric system facilitates a challenge-response type of mechanism, thereby ensuring that the system is interacting with a *live* user. Note that a challenge-response mechanism can be initiated in unibiometric systems also (e.g., system prompts "Please say 1-2-5-7", "Blink twice and move your eyes to the right", "Change your facial expression by smiling", etc.).
4. Multibiometric systems also effectively address the problem of noisy data. When the biometric signal acquired from a single trait is corrupted with noise, the availability of other (less noisy) traits may aid in the reliable determination of identity. Some systems take into account the *quality* of the individual biometric signals during the fusion process. This is especially important when recognition has to take place in adverse conditions where certain biometric traits cannot be reliably extracted. For example, in the presence of ambient acoustic noise, when an individual's voice characteristics cannot be accurately measured, the facial characteristics may be used by the multibiometric system to perform authentication. Estimating the quality of the acquired data is in itself a challenging problem but, when appropriately done, can reap significant benefits in a multibiometric system.
5. These systems also help in the *continuous* monitoring or tracking of an individual in situations when a single trait is not sufficient. Consider a biometric system that uses a 2D camera to procure the face and gait information of a person walking down a crowded aisle. Depending upon the distance and pose of the subject with respect to the camera, both these characteristics may or may not be simultaneously available. Therefore, either (or both) of these traits can be used depending upon the location of the individual with respect to the acquisition system thereby permitting the continuous monitoring of the individual.

6. A multibiometric system may also be viewed as a fault tolerant system which continues to operate even when certain biometric sources become unreliable due to sensor or software malfunction, or deliberate user manipulation. The notion of fault tolerance is especially useful in large-scale authentication systems involving a large number of subjects (such as a border control application).

11.3 Taxonomy of Multibiometric Systems

A multibiometric system relies on the evidence presented by multiple sources of biometric information. Based on the nature of these sources, a multibiometric system can be classified into one of the following six categories [60]: multi-sensor, multi-algorithm, multi-instance, multi-sample, multimodal, or hybrid (see Fig. 11.2).

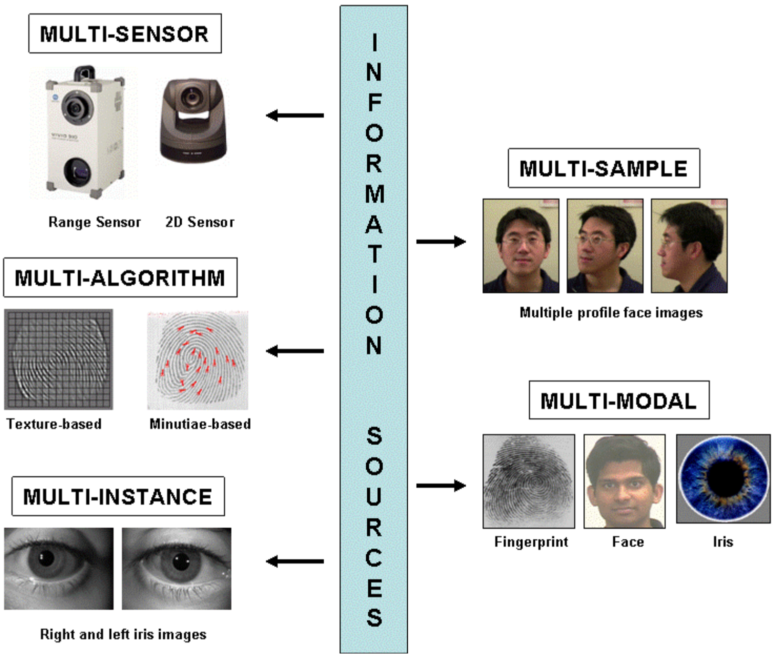

Fig. 11.2 Sources of information for biometric fusion.

1. Multi-sensor systems: Multi-sensor systems employ multiple sensors to capture a single biometric trait of an individual. For example, a face recognition system may deploy multiple 2D cameras to acquire the face image of a subject [37]; an infrared sensor may be used in conjunction with a visible-light sensor to acquire the subsurface information of a person's face [8, 31, 67]; a multispectral camera

may be used to acquire images of the iris, face or finger [48, 62]; or optical and capacitive sensors may be used to image the fingerprint of a subject [40]. The use of multiple sensors, in some instances, can result in the acquisition of complementary information that can enhance the recognition ability of the system. For example, based on the nature of illumination due to ambient lighting, the infrared and visible-light images of a person's face can present different levels of information resulting in enhanced matching accuracy. Similarly, the performance of a 2D face matching system can be improved by utilizing the shape information presented by 3D range images.

2. Multi-algorithm systems: In some cases, invoking multiple feature extraction and/or matching algorithms on the same biometric data can result in improved matching performance. Multi-algorithm systems consolidate the output of multiple feature extraction algorithms, or that of multiple matchers operating on the same feature set. These systems do not necessitate the deployment of new sensors and, hence, are cost-effective compared to other types of multibiometric systems. But on the other hand, the introduction of new feature extraction and matching modules can increase the computational complexity of these systems. Ross et al. [59] describe a fingerprint recognition system that utilizes minutiae as well as texture information to represent and match fingerprint images. Lu et al. [39] discuss a face recognition system that combines three different feature extraction schemes (Principal Component Analysis (PCA), Independent Component Analysis (ICA) and Linear Discriminant Analysis (LDA)).
3. Multi-instance systems: These systems use multiple instances of the same body trait and have also been referred to as multi-unit systems in the literature. For example, the left and right index fingers, or the left and right irises of an individual, may be used to verify an individual's identity [29, 54]. FBI's IAFIS combines the evidence of all ten fingers to determine a matching identity in the database. These systems can be cost-effective if a single sensor is used to acquire the multi-unit data in a sequential fashion, but this can increase data acquisition time thereby causing inconvenience to the user. Thus, in some instances, it may be desirable to obtain the multi-unit data simultaneously thereby demanding the design of an effective (and possibly more expensive) acquisition device.
4. Multi-sample systems: A single sensor may be used to acquire multiple samples of the same biometric trait in order to account for the variations that can occur in the trait, or to obtain a more complete representation of the underlying trait. A face system, for example, may capture (and store) the frontal profile of a person's face along with the left and right profiles in order to account for variations in the facial pose. Similarly, a fingerprint system equipped with a small size sensor may acquire multiple dab prints of an individual's finger in order to obtain images of various regions of the fingerprint. A mosaicing scheme may then be used to stitch the multiple impressions and create a composite image. One of the key issues in a multi-sample system is determining the *number* of samples that have to be acquired from an individual. It is important that the procured samples represent the *variability* as well as the *typicality* of the individual's biometric data. To this end, the desired relationship between the samples has to be established before-hand

in order to optimize the benefits of the integration strategy. For example, a face recognition system utilizing both the frontal- and side-profile images of an individual may stipulate that the side-profile image should be a three-quarter view of the face [22,47]. Alternately, given a set of biometric samples, the system should be able to automatically select the "optimal" subset that would best represent the individual's variability. Uludag et al. [71] discuss two such schemes in the context of fingerprint recognition.

5. Multimodal systems: Multimodal systems establish identity based on the evidence of multiple biometric traits. For example, some of the earliest multimodal biometric systems utilized face and voice features to establish the identity of an individual [4, 6, 11]. Physically uncorrelated traits (e.g., fingerprint and iris) are expected to result in better *improvement* in performance than correlated traits (e.g., voice and lip movement). The cost of deploying these systems is substantially more due to the requirement of new sensors and, consequently, the development of appropriate user interfaces. The identification accuracy can be significantly improved by utilizing an increasing number of traits although the *curse-of-dimensionality* phenomenon would impose a bound on this number. The number of traits used in a specific application will also be restricted by practical considerations such as the cost of deployment, enrollment time, throughput time, expected error rate, user habituation issues, etc.
6. Hybrid systems: Chang et al. [7] use the term *hybrid* to describe systems that integrate a subset of the five scenarios discussed above. For example, Brunelli et al. [6] discuss an arrangement in which two speaker recognition algorithms are combined with three face recognition algorithms at the match score and rank levels via a HyperBF network. Thus, the system is multi-algorithmic as well as multimodal in its design.

11.4 Levels of fusion

Based on the type of information available in a certain module, different levels of fusion may be defined. Sanderson and Paliwal [64] categorize the various levels of fusion into two broad categories: pre-classification or fusion *before* matching, and post-classification or fusion *after* matching (see Figure 11.3). Such a categorization is necessary since the amount of information available for fusion reduces drastically once the matcher has been invoked. Pre-classification fusion schemes typically require the development of new matching techniques (since the matchers used by the individual sources may no longer be relevant) thereby introducing additional challenges. Pre-classification schemes include fusion at the sensor (or raw data) and the feature levels while post-classification schemes include fusion at the match score, rank and decision levels.

1. Sensor-level fusion: The raw biometric data (e.g., a face image) acquired from an individual represents the richest source of information although it is expected to be contaminated by noise (e.g., non-uniform illumination, background clutter,

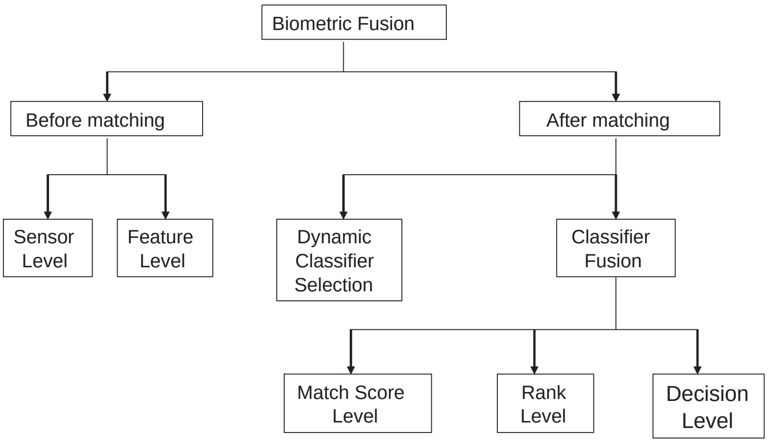

Fig. 11.3 Fusion can be accomplished at various levels in a biometric system.

etc.). Sensor-level fusion refers to the consolidation of (a) raw data obtained using multiple sensors, or (b) multiple snapshots of a biometric using a single sensor [61, 65].

2. Feature-level fusion: In feature-level fusion, the feature sets originating from multiple biometric algorithms are consolidated into a single feature set by the application of appropriate feature normalization, transformation and reduction schemes. The primary benefit of feature-level fusion is the detection of correlated feature values generated by different biometric algorithms and, in the process, identifying a salient set of features that can improve recognition accuracy. Eliciting this feature set typically requires the use of dimensionality reduction methods and, therefore, feature-level fusion assumes the availability of a large number of training data. Also, the feature sets being fused are typically expected to reside in commensurate vector space in order to permit the application of a suitable matching technique upon consolidating the feature sets [58, 68].
3. Score-level fusion: In score-level fusion the match scores output by multiple biometric matchers are combined to generate a new match score (a scalar) that can be subsequently used by the verification or identification modules for rendering an identity decision. Fusion at this level is the most commonly discussed approach in the biometric literature primarily due to the ease of accessing and processing match scores (compared to the raw biometric data or the feature set extracted from the data). Fusion methods at this level can be broadly classified into three categories [60]: density-based schemes [12, 70], transformation-based schemes [26] and classifier-based schemes [72].
4. Rank-level fusion: When a biometric system operates in the identification mode, the output of the system can be viewed as a ranking of the enrolled identities. In this case, the output indicates the set of possible matching identities sorted in decreasing order of confidence. The goal of rank-level fusion schemes is to

consolidate the ranks output by the individual biometric subsystems in order to derive a consensus rank for each identity. Ranks provide more insight into the decision-making process of the matcher compared to just the identity of the best match, but they reveal less information than match scores. However, unlike match scores, the rankings output by multiple biometric systems are comparable. As a result, no normalization is needed and this makes rank-level fusion schemes simpler to implement compared to the score-level fusion techniques [23].

5. Decision-level fusion: Many commercial off-the-shelf (COTS) biometric matchers provide access only to the final recognition decision. When such COTS matchers are used to build a multibiometric system, only decision-level fusion is feasible. Methods proposed in the literature for decision-level fusion include "AND" and "OR" rules [13], majority voting [36], weighted majority voting [34], Bayesian decision fusion [74], the Dempster-Shafer theory of evidence [74], and behavior knowledge space [24].

11.5 Incorporating Ancillary Information

Another category of multibiometric systems combine primary biometric identifiers (such as face and fingerprint) with soft biometric attributes (such as gender, height, weight, eye color, etc.). Soft biometric traits cannot be used to distinguish individuals reliably since the same attribute is likely to be shared by several different people in the target population. However, when used in conjunction with primary biometric traits, the performance of the authentication system can be significantly enhanced [25]. Soft biometric attributes also help in filtering (or indexing) large biometric databases by limiting the number of entries to be searched in the database. For example, if it is determined (automatically or manually) that the subject is an "Asian Male", then the system can constrain its search to only those identities in the database labeled with these attributes. Alternately, soft biometric traits can be used in surveillance applications to decide if at all primary biometric information has to be acquired from a certain individual. Automated techniques to estimate soft biometric characteristics is an ongoing area of research and is likely to benefit law enforcement and border control biometric applications.

Some biometric systems incorporate data quality into the fusion process. This strategy is referred to as *quality-based fusion*. The purpose is (a) to automatically assign weights to the participating modalities thereby mitigating the errors introduced by poor quality input data [45], or (b) to appropriately invoke the modalities in a cascade fashion thereby maximizing recognition accuracy [16]. Soft biometric data and quality indices are referred to as ancillary information in the context of biometric fusion.

In quality-based fusion (e.g., [5, 18, 30, 32, 45]), the quality associated with the template (i.e., gallery) as well as the query (i.e., probe) biometric sample are taken into account. For assessing the quality of a biometric sample, a number of measures have been proposed in the literature (e.g., fingerprint [10, 19], iris [9], face [20], speech [46], signature [44], and classifier-dependent confidence measures [3, 50]). These quality measures, in general, aim to quantify the degree of excellence or conformance of biometric samples to some predefined criteria known to influence the performance of the system.

Depending on their role, there are at least two ways in which quality measures can be incorporated into a fusion classifier - either as a control parameter (primary role) or as evidence (secondary role). In their primary role, quality measures are used to modify the way a fusion classifier is trained or tested. For example, quality measures have been incorporated into the following classifiers: Bayesian-based classifier [5], reduced polynomial classifier [69], support vector machine [18], and fixed-rule fusion [17]. In their secondary role, quality measures are often concatenated with the outputs of individual matchers (such as match scores) and the ensuing 'feature' is input to a fusion classifier as discussed in [30] (logistic regression classifier) and [45] (Bayesian classifier). The use of Bayesian networks to gauge the complex relationship between expert outputs and quality measures (e.g., [42]) has also been explored. The work in [52] takes into account an array of quality measures rather than summarizing the quality as a scalar value. By means of grouping the multi faceted quality measures into multiple clusters, a unique fusion strategy can be devised for each cluster.

Quality measures have also been used to improve biometric device interoperability [1, 51]. Such an approach is commonly used in speaker verification where different strategies are invoked based on the microphone type [2].

The notion of quality is closely related to that of *reliability*. In [33], the estimated reliability for each biometric modality was used for combining symbolic-level decisions and in [38, 50, 57, 73], score-level fusion was considered. In [38, 57, 73], the term "failure prediction" is used instead of "reliability". Such information derived solely from the output of individual matchers (rather than explicit quality measures) has been demonstrated to be effective for (a) single biometric modalities, (b) fusion across multiple sensors for a single biometric modality, and (c) fusion across different machine learning techniques. In [50], the notion of reliability was represented by the *margin*, a concept used in large-margin classifiers [66]. However, a precise definition for reliability and the procedure used for estimating it are open research issues and further work is essential in this regard.

11.6 Benefits of Combining Multiple Biometric Experts: Three Case Studies

The literature on biometrics is replete with examples demonstrating the benefits of multibiometric fusion.[1] In this section, three examples illustrating the benefits of multibiometric systems as well as the design issues involved are presented. All

[1] The term "expert" is sometimes used to refer to individual biometric matchers or modalities used in a multibiometric system.

the examples are based on results reported in the literature on public datasets. The first case illustrates the potential of multibiometric systems to improve matching accuracy. The second case illustrates the benefit of using quality measures in the fusion scheme. The third case demonstrates the possibility of optimizing the cost of authentication for a given target accuracy by managing the choice of biometric traits and matchers. Such a cost-based analysis enables one to decide, for instance, if combining multiple biometric traits is better than combining multiple samples of the same biometric (by reusing the same device), as the latter is less expensive.

11.6.1 On the Complementarity of Multimodal Experts

The first case study which illustrates the merits of both multimodal and multi-algorithm fusion [63] involves fusing three different modalities: face, voice and lip movement. Two face recognition algorithms, two speaker recognition algorithms and one lip dynamic recognition algorithm were used. This multibiometric system was evaluated on the XM2VTS data base [43] producing match scores according to the Lausanne Experimental Protocol in Configuration I [43] (see Table 11.1). Although the performances of the individual experts were unremarkable (with the exception of one speaker recognition algorithm), the fusion of these biometric experts by simple weighted averaging resulted in improved performance, as summarized in Table 11.2. The results show that multimodal fusion has the potential to ameliorate the performance of the single best expert even if some of the individual recognition algorithms have error rates that are an order of magnitude worse than the best expert. Interestingly, the combination of the best experts from each of the three modalities is only marginally better than the best performing speaker recognition algorithm. In the second row of Table 11.2, it is observed that the weights assigned to the weaker algorithms are greater than those associated with the best algorithm of each modality. It appears that the diversity offered by these weaker algorithms has led to an improved matching accuracy after fusion.

Algorithm	Threshold	FRR	FAR
Lips	0.50	14.00 %	12.67 %
Face 1	0.21	5.00 %	4.45 %
Face 2	0.50	6.00 %	8.12 %
Voice 1	0.50	7.00 %	1.42 %
Voice 2	0.50	0.00 %	1.48 %

Table 11.1 Performance of multiple biometric algorithms (i.e., experts) on the test set (Configuration I). Results are from the work reported in [63].

Experts	Weights	Threshold	FRR	FAR
Lips, Face 1, Voice 2	0.27, 0.23, 0.50	0.51	0.00 %	1.31 %
Face 1, Face 2, Voice 1, Voice 2	0.02, 0.06, 0.87, 0.05	0.50	0.00 %	0.52 %
Lips, Face 1, Face 2, Voice 1, Voice 2	0.03, 0.01, 0.04, 0.89, 0.03	0.50	0.00 %	0.29%

Table 11.2 Fusion using the simple sum rule (Configuration I). Results are from the work reported in [63].

11.6.2 Benefits of Quality-based Fusion

The second example demonstrates the benefits of using quality measures in fusion [30]. In the referenced study, logistic regression was used as a fusion classifier, producing an output that approximates the posterior probability of observing a vector of match scores (denoted as x); the elements of the vector correspond to the match scores generated by the individual experts.

Quality measures (denoted as q) are then considered as additional inputs to the fusion module. The interaction of quality measures and match scores is explicitly modeled by feeding three variants of input to the fusion classifier: $[x,q]$ (i.e., augmenting x and q via concatenation), $[x,x \otimes q]$ (where $\otimes$ denotes a tensor product) and $[x,q,x \otimes q]$. If there are N_q terms in q and N_x terms in x, the tensor product between q and x produces $N_q \times N_x$ elements, hence, providing the fusion classifier with an additional degree of freedom to model the pair-wise product elements generated by the two vectors. Thus, the input to the fusion module has four possible arrangements: x, $[x,q]$, $[x,x \otimes q]$ and $[x,q,x \otimes q]$.

The results of using these input arrangements on the XM2VTS database with both standard and degraded data are shown in Figure 11.4. Six face recognition algorithms and one speaker recognition algorithm were used in the experiment. Since the voice modality was included in all fusion tasks, the number of ways in which the experts can be combined is $2^6 - 1 = 63$. Each bar in this figure represents a statistic measuring the relative difference in error between a fusion system that does *not* use quality and one of the arrangements mentioned above. As can be observed, the use of quality measures can reduce the verification error (measured using the Equal Error Rate (EER)) by as much as 40%. Such an improvement is possible especially when both the face and voice biometric modalities contain significantly different quality types. In the experimental setting, the speech was artificially corrupted by uniformly distributed additive noise of varying magnitude (from 0dB to 20dB) whereas the face data contained either well illuminated or side illuminated face images. More research is needed to account for cases when the noise characteristics are ill-defined or unknown.

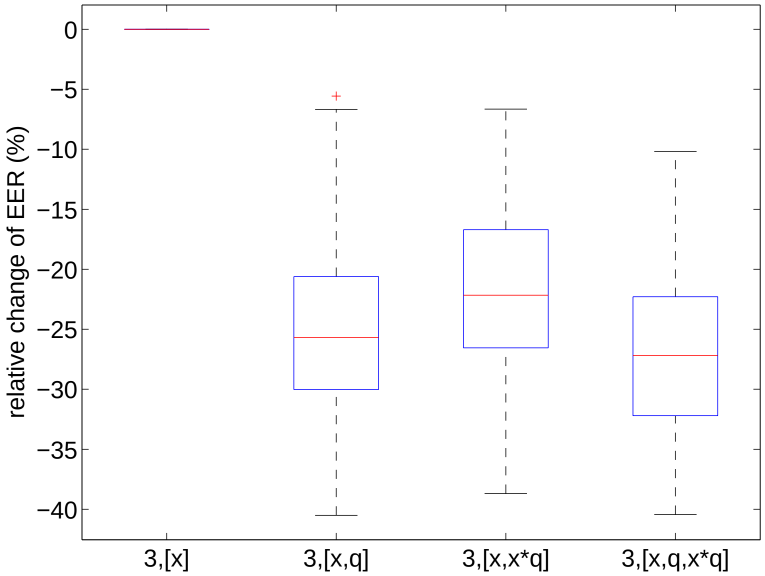

Fig. 11.4 Relative change in *a posteriori EER* (%) when combining the outputs of multiple biometric algorithms on the "good" and "degraded" subsets of the XM2VTS face database according to the modified Lausanne Protocol Configuration I [30]. Each bar shows the distribution of 63 values corresponding to the 63 possible face and speaker fusion tasks. The first and third quantiles are depicted by a bounding box and the median value by a horizontal red bar. In arrangement $[x]$, the quality information is not used. The mean absolute performance of the first bar is about 3 percent (EER) whereas that of the remaining three quality-based fusion systems is about 2 percent.

11.6.3 Multibiometric Expert Selection

The third example casts the fusion problem as an optimization problem. Since using more than one biometric modality often implies the need for additional hardware and software resources as well as authentication time (thus inconveniencing the user), it is reasonable to attribute an abstract cost to each additional biometric device. The goal of optimization in this context can be formulated as follows: determine the subset of candidate biometric experts that minimizes the overall cost of operation while exhibiting reasonable matching accuracy. Ideally, such a cost-sensitive optimization criterion should be robust to population mismatch, where one attempts to design a fusion classifier based on a development set of subjects and apply the resultant classifier to a target set of subjects not present in the development set. One possible objective function (criterion) for this optimization problem could be the error rate of the fusion classifier.

An example of a cost-sensitive performance curve for a verification experiment is shown in Figure 11.5. This experiment was conducted on the Biosecure DS2 database. [2] Match scores were generated using 8 different recognition algorithms pertaining to the face, fingerprint and iris modalities. In order to carry out the experiment, the match score database was partitioned into two: the development set, for training, and the evaluation set, for testing. The genuine and impostor user populations were different for these two sets. Two theoretical error measures (the Chernoff and Bhattacharyya bounds) and two empirical EER measures (corresponding to two Bayesian classifiers) were used as objective functions for the optimization problem. The first Bayesian classifier was designed by estimating the genuine and impostor densities using a Gaussian mixture model and using the likelihood ratio test statistic to classify an arbitrary score (GMM Classifier). The second classifier was designed by estimating a decision boundary to separate the genuine and impostor scores; the boundary itself was estimated using Quadratic Discriminant Analysis (QDA Classifier). See [53] for details. For estimating the empirical error rate on the development set, a two-fold cross-validation procedure was employed. The average EER of the two folds was used as an indicator of the error on the development set. The classifier trained using the development set was then used to assess the empirical error rate on the evaluation set in a similar way. In order to compute the theoretical error bounds, the genuine and impostor match scores were modeled using normal distributions whose parameters were estimated using the development/evaluation data set.

These error measures were computed for both the development and evaluation sets for various combinations of experts. The cost of each combination was also computed. In the fusion problem considered here (see Figure 11.5), the cost ranges from 1 (using a single expert) to 4.5 (using all 8 experts). Costs were assigned as follows: the use of one system is charged a unit cost; subsequent reuse of the system (e.g., multiple fingers in a multi-unit system) is charged 0.3 units. Thus, using a face, an iris and 2 fingers will incur a cost of $1+1+1+0.3 = 3.3$ units.

When eight experts (i.e., face, one iris and six fingers) are used, the search space has $2^8 - 1 = 255$ elements. We plot here a “rank-one” *cost-sensitive* performance curve (performance versus cost) where the performance has been computed on the evaluation set. Since the goal is to achieve minimum error with minimum cost, a curve towards the lower left corner is the ideal one. This curve is called a rank-one curve because it uses the development set to determine the best performing combination-of-experts at each cost value. Similarly, a rank-two cost-sensitive curve would be computed by using the development set to determine the top two performing combination-of-experts at each cost value, and then reporting the best of these two combinations on the evaluation set. With enough rank order, the performance curve will tend to the oracle (the ideal curve with error estimated on the evaluation set). While the rank-one curve using the Bhattacharyya bound as the objective function is satisfactory, the rank-three curve was observed to exhibit exactly the same characteristics as the oracle for the QDA classifier and the rank-five curve was observed to exhibit exactly the same characteristics as the oracle for the GMM classifier. When the empirical error rate was used as the objective function, a rank-six curve was needed to achieve the performance of the oracle for the QDA classifier

[2] The data set is available for download at http://face.ee.surrey.ac.uk/qfusion.

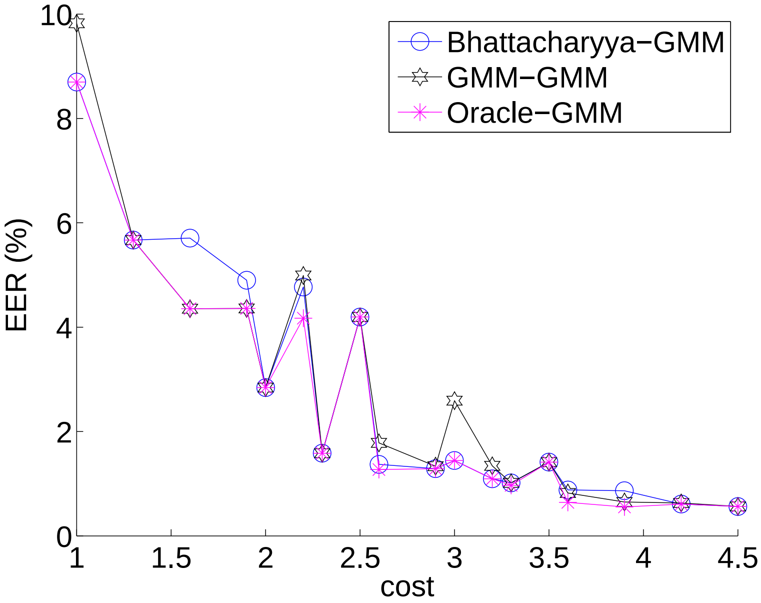

Fig. 11.5 The rank-one cost-sensitive performance curve using the GMM Bayesian classifier on all 255 combinations in the Biosecure DS2 fusion database. The legend "Bhattacharyya-GMM" refers to optimization using the Bhattacharyya criterion on the development set and measuring the performance of the GMM Bayesian classifier on the evaluation set. The curve "GMM-GMM" should be interpreted similarly. The Oracle-GMM is the performance of the GMM Bayesian classifier on the evaluation set. The experimental results using the QDA Bayesian classifier are similar to this figure (not shown here).

while more than rank-ten was needed for achieving the oracle for the GMM classifier.

11.7 Open Issues and Challenges

In this section, we will highlight some of the issues in multibiometrics which call for further research.

- **Fusion Architecture:** The range of possible fusion configurations encompassing serial (cascade), parallel and hybrid modes of operation is very large. While the parallel fusion strategy is most commonly used, there are additional advantages in exploring serial fusion where the experts are considered one at a time. It offers the possibility of making reliable decisions with a few experts, leaving only difficult problems to be handled by the remaining (possibly more expensive) experts. However, automatically deducing these configurations is an open problem and requires more research.

- **Correlated Experts:** An important consideration when adopting a fusion strategy is to model the statistical dependency among the expert outputs. For instance, in a multi-algorithm setting, several experts may rely on the same biometric sample and so higher dependency is expected among the expert outputs. On the other hand, in a multimodal setting, the pool of experts is likely to be statistically independent. Assessing the impact of correlated experts on overall matching accuracy is an interesting problem that has received very little attention in the biometric literature (see [30]).
- **Expert selection:** Expert selection can be cast as a feature selection problem, as illustrated in [49]. However, directly applying such a technique to biometric authentication is difficult for several reasons. In Section 11.6, for instance, we have seen that the optimal set of experts determined using a development population of users may not be the best for the target set of users. Further, the "Doddington Zoo" effect would result in an asymmetrical distribution of errors across the user population [14]. Another issue is raised by cost considerations. Conciliating both the operational cost and matching accuracy into a single criterion is a difficult task.

  A related problem, known as dynamic expert selection, arises in the context of serial fusion. In dynamic expert selection, a fusion classifier may decide which expert would be the *most informative* even before the biometric data is acquired from the user. In the recent Multimodal Biometric benchmark evaluation organized by the Biosecure (EU-funded) project, the use of a dynamic fusion strategy proved to be very promising in achieving good performance while minimizing costs.

Other topics of research in multibiometrics include (a) protecting multibiometric templates; (b) indexing multimodal databases; (c) consolidating biometric sources in highly unconstrained environments; (d) designing dynamic fusion algorithms to address the problem of incomplete input data; (e) predicting the matching performance of a multibiometric system; and (f) continuous monitoring of an individual using multiple traits.

11.8 Summary

Multibiometric systems are expected to enhance the recognition accuracy of a personal authentication system by reconciling the evidence presented by multiple sources of information. In this chapter, the different sources of biometric information as well as the type of information that can be consolidated was presented. Typically, early integration strategies (e.g., feature-level) are expected to result in better performance than late integration (e.g., score-level) strategies. However, it is difficult to predict the performance gain due to each of these strategies prior to invoking the fusion methodology. The use of ancillary information, such as soft biometrics and data quality, can further improve the performance of a multibiometric system if an appropriate fusion strategy is used. The three case studies discussed in this chapter highlight the benefits of fusion.

While the *availability* of multiple sources of biometric information (pertaining either to a single trait or to multiple traits) may present a compelling case for fusion, the *correlation* between the sources has to be examined before determining their suitability for fusion. Combining uncorrelated or negatively correlated sources is expected to result in a better improvement in matching performance than combining positively correlated sources [35]. However, defining an appropriate diversity measure to predict fusion performance has been elusive. Thus, there are several open challenges in the multibiometric field that require further research. Nevertheless, it is becoming increasingly apparent that multibiometric systems will have a profound impact on how identity is established in the 21st century.

Acknowledgements Norman Poh was supported partially by the advanced researcher fellowship PA0022_121477 of the Swiss National Science Foundation, by the EU-funded Mobio project grant IST-214324 and by the Engineering and Physical Sciences Research Council (EPSRC) Research Grant GR/S46543. Arun Ross was supported by US NSF CAREER grant number IIS 0642554.

References

1. F. Alonso-Fernandez, J. Fierrez, D. Ramos, and J. Ortega-Garcia. Dealing with sensor interoperability in multi-biometrics: The upm experience at the biosecure multimodal evaluation 2007. In *Proc. of SPIE Defense and Security Symposium, Workshop on Biometric Technology for Human Identification*, 2008.
2. R. Auckenthaler, M. Carey, and H. Lloyd-Thomas. Score Normalization for Text-Independant Speaker Verification Systems. *Digital Signal Processing (DSP) Journal*, 10:42–54, 2000.
3. S. Bengio, C. Marcel, S. Marcel, and J. Marithoz. Confidence Measures for Multimodal Identity Verification. *Information Fusion*, 3(4):267–276, 2002.
4. E. S. Bigun, J. Bigun, B. Duc, and S. Fischer. Expert Conciliation for Multimodal Person Authentication Systems using Bayesian Statistics. In *First International Conference on Audio- and Video-based Biometric Person Authentication (AVBPA)*, pages 291–300, Crans-Montana, Switzerland, March 1997.
5. J. Bigun, J. Fierrez-Aguilar, J. Ortega-Garcia, and J. Gonzalez-Rodriguez. Multimodal Biometric Authentication using Quality Signals in Mobile Communnications. In *12th Int'l Conf. on Image Analysis and Processing*, pages 2–13, Mantova, 2003.
6. R. Brunelli and D. Falavigna. Person Identification Using Multiple Cues. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 17(10):955–966, October 1995.
7. K. I. Chang, K. W. Bowyer, and P. J. Flynn. An Evaluation of Multimodal 2D+3D Face Biometrics. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 27(4):619–624, April 2005.
8. X. Chen, P. J. Flynn, and K. W. Bowyer. IR and Visible Light Face Recognition. *Computer Vision and Image Understanding*, 99(3):332–358, September 2005.
9. Y. Chen, S. Dass, and A. Jain. Localized iris image quality using 2-d wavelets. In *Proc. Int'l Conf. on Biometrics (ICB)*, pages 373–381, Hong Kong, 2006.
10. Y. Chen, S.C. Dass, and A.K. Jain. Fingerprint Quality Indices for Predicting Authentication Performance. In *LNCS 3546, 5th Int'l. Conf. Audio- and Video-Based Biometric Person Authentication (AVBPA 2005)*, pages 160–170, New York, 2005.

11. C. C. Chibelushi, J. S. D. Mason, and F. Deravi. Feature-level Data Fusion for Bimodal Person Recognition. In *Proceedings of the Sixth International Conference on Image Processing and Its Applications*, volume 1, pages 399–403, Dublin, Ireland, July 1997.
12. S. C. Dass, K. Nandakumar, and A. K. Jain. A Principled Approach to Score Level Fusion in Multimodal Biometric Systems. In *Proceedings of Fifth International Conference on Audio- and Video-based Biometric Person Authentication (AVBPA)*, pages 1049–1058, Rye Brook, USA, July 2005.
13. J. Daugman. Combining Multiple Biometrics. Available at http://www.cl.cam.ac.uk/users/jgd1000, 2000.
14. G. Doddington, W. Liggett, A. Martin, M. Przybocki, and D. Reynolds. Sheep, Goats, Lambs and Wolves: A Statistical Analysis of Speaker Performance in the NIST 1998 Speaker Recognition Evaluation. In *Int'l Conf. Spoken Language Processing (ICSLP)*, Sydney, 1998.
15. A. Eriksson and P. Wretling. How Flexible is the Human Voice? A Case Study of Mimicry. In *Proceedings of the European Conference on Speech Technology*, pages 1043–1046, Rhodes, 1997.
16. E. Erzin, Y. Yemez, and A. M. Tekalp. Multimodal Speaker Identification Using an Adaptive Classifier Cascade Based on Modality Reliability. *IEEE Transactions on Multimedia*, 7(5):840–852, October 2005.
17. O. Fatukasi, J. Kittler, and N. Poh. Quality Controlled Multimodal Fusion of Biometric Experts. In *12th Iberoamerican Congress on Pattern Recognition CIARP*, pages 881–890, Via del Mar-Valparaiso, Chile, 2007.
18. J. Fierrez-Aguilar, J. Ortega-Garcia, J. Gonzalez-Rodriguez, and J. Bigun. Kernel-Based Multimodal Biometric Verification Using Quality Signals. In *Proc. of SPIE Defense and Security Symposium, Workshop on Biometric Technology for Human Identification*, volume 5404, pages 544–554, 2004.
19. H. Fronthaler, K. Kollreider, J. Bigun, J. Fierrez, F. Alonso-Fernandez, J. Ortega-Garcia, and J. Gonzalez-Rodriguez. Fingerprint image-quality estimation and its application to multialgorithm verification. *IEEE Trans. on Information Forensics and Security*, 3:331–338, 2008.
20. X. Gao, R. Liu, S. Z. Li, and P. Zhang. Standardization of face image sample quality. In *LNCS 4642, Proc. Int'l Conf. Biometrics (ICB'07)*, pages 242–251, Seoul, 2007.
21. W. R. Harrison. *Suspect Documents, their Scientific Examination*. Nelson-Hall Publishers, 1981.
22. H. Hill, P. G. Schyns, and S. Akamatsu. Information and Viewpoint Dependence in Face Recognition. *Cognition*, 62(2):201–222, February 1997.
23. T. K. Ho, J. J. Hull, and S. N. Srihari. Decision Combination in Multiple Classifier Systems. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 16(1):66–75, January 1994.
24. Y. S. Huang and C. Y. Suen. Method of Combining Multiple Experts for the Recognition of Unconstrained Handwritten Numerals. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 17(1):90–94, January 1995.
25. A. K. Jain, K. Nandakumar, X. Lu, and U. Park. Integrating Faces, Fingerprints and Soft Biometric Traits for User Recognition. In *Proceedings of ECCV International Workshop on Biometric Authentication (BioAW)*, volume LNCS 3087, pages 259–269, Prague, Czech Republic, May 2004. Springer.
26. A. K. Jain, K. Nandakumar, and A. Ross. Score Normalization in Multimodal Biometric Systems. *Pattern Recognition*, 38(12):2270–2285, December 2005.
27. A. K. Jain and A. Ross. Multibiometric Systems. *Communications of the ACM, Special Issue on Multimodal Interfaces*, 47(1):34–40, January 2004.
28. A. K. Jain, A. Ross, and S. Prabhakar. An Introduction to Biometric Recognition. *IEEE Transactions on Circuits and Systems for Video Technology, Special Issue on Image- and Video-Based Biometrics*, 14(1):4–20, January 2004.
29. J. Jang, K. R. Park, J. Son, and Y. Lee. Multi-unit Iris Recognition System by Image Check Algorithm. In *Proceedings of International Conference on Biometric Authentication (ICBA)*, pages 450–457, Hong Kong, July 2004.

30. J. Kittler, N. Poh, O. Fatukasi, K. Messer, K. Kryszczuk, J. Richiardi, and A. Drygajlo. Quality Dependent Fusion of Intramodal and Multimodal Biometric Experts. In *Proc. of SPIE Defense and Security Symposium, Workshop on Biometric Technology for Human Identification*, volume 6539, 2007.
31. A. Kong, J. Heo, B. Abidi, J. Paik, and M. Abidi. Recent Advances in Visual and Infrared Face Recognition - A Review. *Computer Vision and Image Understanding*, 97(1):103–135, January 2005.
32. K. Kryszczuk and A. Drygajlo. Credence estimation and error prediction in biometric identity verification. *Signal Processing*, 88:916–925, 2008.
33. Krzysztof Kryszczuk, Jonas Richiardi, Plamen Prodanov, and Andrzej Drygajlo. Reliability-based decision fusion in multimodal biometric verification systems. *EURASIP Journal of Advances in Signal Processing*, 2007.
34. L. I. Kuncheva. *Combining Pattern Classifiers - Methods and Algorithms*. Wiley, 2004.
35. L. I. Kuncheva, C. J. Whitaker, C. A. Shipp, and R. P. W. Duin. Is Independence Good for Combining Classifiers? In *Proceedings of International Conference on Pattern Recognition (ICPR)*, volume 2, pages 168–171, Barcelona, Spain, 2000.
36. L. Lam and C. Y. Suen. Application of Majority Voting to Pattern Recognition: An Analysis of its Behavior and Performance. *IEEE Transactions on Systems, Man, and Cybernetics, Part A: Systems and Humans*, 27(5):553–568, 1997.
37. J. Lee, B. Moghaddam, H. Pfister, and R. Machiraju. Finding Optimal Views for 3D Face Shape Modeling. In *Proceedings of the IEEE International Conference on Automatic Face and Gesture Recognition (FG)*, pages 31–36, Seoul, Korea, May 2004.
38. W. Li, X. Gao, and T.E. Boult. Predicting biometric system failure. *Proceedings of the IEEE International Conference on Computational Intelligence for Homeland Security and Personal Safety*, pages 57–64, 31 2005-April 1 2005.
39. X. Lu, Y. Wang, and A. K. Jain. Combining Classifiers for Face Recognition. In *IEEE International Conference on Multimedia and Expo (ICME)*, volume 3, pages 13–16, Baltimore, USA, July 2003.
40. G. L. Marcialis and F. Roli. Fingerprint Verification by Fusion of Optical and Capacitive Sensors. *Pattern Recognition Letters*, 25(11):1315–1322, August 2004.
41. T. Matsumoto, H. Matsumoto, K. Yamada, and S. Hoshino. Impact of Artificial Gummy Fingers on Fingerprint Systems. In *Optical Security and Counterfeit Deterrence Techniques IV, Proceedings of SPIE*, volume 4677, pages 275–289, San Jose, USA, January 2002.
42. D. E. Maurer and J. P. Baker. Fusing multimodal biometrics with quality estimates via a bayesian belief network. *Pattern Recognition*, 41(3):821–832, 2007.
43. K Messer, J Matas, J Kittler, J Luettin, and G Maitre. Xm2vtsdb: The extended m2vts database. In *Second International Conference on Audio and Video-based Biometric Person Authentication*, 1999.
44. S. Muller and O. Henniger. Evaluating the biometric sample quality of handwritten signatures. In *LNCS 3832, Proc. Int'l Conf. Biometrics (ICB'07)*, pages 407–414, 2007.
45. K. Nandakumar, Y. Chen, S. C. Dass, and A. K. Jain. Likelihood ratio based biometric score fusion. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 30:342–347, 2008.
46. National Institute of Standards and Technology. Nist speech quality assurance package 2.3 documentation.
47. A. O'Toole, H. Bulthoff, N. Troje, and T. Vetter. Face Recognition across Large Viewpoint Changes. In *Proceedings of the International Workshop on Automatic Face- and Gesture-Recognition (IWAFGR)*, pages 326–331, Zurich, Switzerland, June 1995.
48. Z. Pan, G. Healey, M. Prasad, and B. Tromberg. Face Recognition in Hyperspectral Images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25(12):1552–1560, December 2003.
49. N. Poh and S. Bengio. Towards Predicting Optimal Subsets of Base-Experts in Biometric Authentication Task. In *LNCS 3361, 1st Joint AMI/PASCAL/IM2/M4 Workshop on Multimodal Interaction and Related Machine Learning Algorithms MLMI*, pages 159–172, Martigny, 2004.

50. N. Poh and S. Bengio. Improving Fusion with Margin-Derived Confidence in Biometric Authentication Tasks. In *LNCS 3546, 5th Int'l. Conf. Audio- and Video-Based Biometric Person Authentication (AVBPA 2005)*, pages 474–483, New York, 2005.
51. N. Poh, T. Bourlai, and J. Kittler. Improving Biometric Device Interoperability by Likelihood Ratio-based Quality Dependent Score Normalization. In *IEEE Conf. on Biometrics: Theory, Applications and Systems*, pages 1–5, Washington, D.C., 2007.
52. N. Poh, G. Heusch, and J. Kittler. On Combination of Face Authentication Experts by a Mixture of Quality Dependent Fusion Classifiers. In *LNCS 4472, Multiple Classifiers System (MCS)*, pages 344–356, Prague, 2007.
53. N. Poh and J. Kittler. On Using Error Bounds to Optimize Cost-sensitive Multimodal Biometric Authentication. In *Proc. 19th Int'l Conf. Pattern Recognition (ICPR)*, 2008.
54. S. Prabhakar and A. K. Jain. Decision-level Fusion in Fingerprint Verification. Technical Report MSU-CSE-00-24, Michigan State University, October 2000.
55. T. Putte and J. Keuning. Biometrical Fingerprint Recognition: Don't Get Your Fingers Burned. In *Proceedings of IFIP TC8/WG8.8 Fourth Working Conference on Smart Card Research and Advanced Applications*, pages 289–303, 2000.
56. N. K. Ratha, J. H. Connell, and R. M. Bolle. An Analysis of Minutiae Matching Strength. In *Proceedings of Third International Conference on Audio- and Video-Based Biometric Person Authentication (AVBPA)*, pages 223–228, Halmstad, Sweden, June 2001.
57. T. P. Riopka and T. E. Boult. Classification enhancement via biometric pattern perturbation. In *Proc. of Audio- and Video-Based Biometric Person Authentication (AVBPA)*, pages 850–859, 2005.
58. A. Ross and R. Govindarajan. Feature Level Fusion Using Hand and Face Biometrics. In *Proceedings of SPIE Conference on Biometric Technology for Human Identification II*, volume 5779, pages 196–204, Orlando, USA, March 2005.
59. A. Ross, A. K. Jain, and J. Reisman. A Hybrid Fingerprint Matcher. *Pattern Recognition*, 36(7):1661–1673, July 2003.
60. A. Ross, K. Nandakumar, and A. K. Jain. *Handbook of Multibiometrics*. Springer, New York, USA, 1st edition, 2006.
61. A. Ross, S. Shah, and J. Shah. Image Versus Feature Mosaicing: A Case Study in Fingerprints. In *Proceedings of SPIE Conference on Biometric Technology for Human Identification III*, pages 620208–1 – 620208–12, Orlando, USA, April 2006.
62. R. K. Rowe and K. A. Nixon. Fingerprint Enhancement Using a Multispectral Sensor. In *Proceedings of SPIE Conference on Biometric Technology for Human Identification II*, volume 5779, pages 81–93, March 2005.
63. U.R. Sanchez and J. Kittler. Fusion of talking face biometric modalities for personal identity verification. In *IEEE Int'l Conf. Acoustics, Speech, and Signal Processing*, volume 5, pages V–V, 2006.
64. C. Sanderson and K. K. Paliwal. Information Fusion and Person Verification Using Speech and Face Information. Research Paper IDIAP-RR 02-33, IDIAP, September 2002.
65. R. Singh, M. Vatsa, A. Ross, and A. Noore. Performance Enhancement of 2D Face Recognition via Mosaicing. In *Proceedings of the 4th IEEE Workshop on Automatic Identification Advanced Technologies (AuotID)*, pages 63–68, Buffalo, USA, October 2005.
66. A. J. Smola and P. J. Bartlett, editors. *Advances in Large Margin Classifiers*. MIT Press, Cambridge, MA, 2000.
67. D. A. Socolinsky, A. Selinger, and J. D. Neuheisel. Face Recognition with Visible and Thermal Infrared Imagery. *Computer Vision and Image Understanding*, 91(1-2):72–114, July-August 2003.
68. B. Son and Y. Lee. Biometric Authentication System Using Reduced Joint Feature Vector of Iris and Face. In *Proceedings of Fifth International Conference on Audio- and Video-Based Biometric Person Authentication (AVBPA)*, pages 513–522, Rye Brook, USA, July 2005.
69. K-A. Toh, W-Y. Yau, E. Lim, L. Chen, and C-H. Ng. Fusion of Auxiliary Information for Multimodal Biometric Authentication. In *LNCS 3072, Int'l Conf. on Biometric Authentication (ICBA)*, pages 678–685, Hong Kong, 2004.

70. B. Ulery, A. Hicklin, C. Watson, W. Fellner, and P. Hallinan. Studies of Biometric Fusion. Technical Report NISTIR 7346, NIST, September 2006.
71. U. Uludag, A. Ross, and A. K. Jain. Biometric Template Selection and Update: A Case Study in Fingerprints. *Pattern Recognition*, 37(7):1533–1542, July 2004.
72. P. Verlinde and G. Cholet. Comparing Decision Fusion Paradigms using k-NN based Classifiers, Decision Trees and Logistic Regression in a Multi-modal Identity Verification Application. In *Proceedings of Second International Conference on Audio- and Video-Based Biometric Person Authentication (AVBPA)*, pages 188–193, Washington D.C., USA, March 1999.
73. B. Xie, T. Boult, V. Ramesh, and Y. Zhu. Multi-camera face recognition by reliability-based selection. *Proceedings of the IEEE International Conference on Computational Intelligence for Homeland Security and Personal Safety*, pages 18–23, Oct. 2006.
74. L. Xu, A. Krzyzak, and C. Y. Suen. Methods for Combining Multiple Classifiers and their Applications to Handwriting Recognition. *IEEE Transactions on Systems, Man, and Cybernetics*, 22(3):418–435, 1992.