Multibiometric Authentication

- An Overview of Recent Developments -

Term Project CS574
Spring 2003
San Diego State University

Uwe M. Bubeck

uwe@ub-net.de

Contents

ABSTRACT

This paper reports on recent research in the area of multibiometric authentication. After outlining the motivation behind this extension of conventional biometrics to incorporate multiple biometric identifiers, the main part of the paper gives a comparative overview of the different architectures of multibiometric systems. In the last part of the paper, I discuss the effects of multibiometrics on the user, before I come to a final conclusion.

1. INTRODUCTION

In recent years, biometric authentication has seen considerable improvements in reliability and accuracy, with some biometrics offering reasonably good overall performance (see [1] for a comparative survey of state-of-the-art biometric authentication technology). However, even the best biometrics to date are still facing numerous problems, some of them inherent to the technology itself. In particular, biometric authentication systems generally suffer from enrollment problems due to non-universal biometric traits, susceptibility to biometric spoofing or insufficient accuracy caused by noisy data acquisition in certain environments.

Multibiometrics are a relatively new approach to overcome those problems. Driven by lower hardware costs, a multibiometric system uses multiple sensors for data acquisition. This allows it to capture multiple samples of a single biometric trait (called multi-sample biometrics) and/or samples of multiple biometric traits (called multi-source or multi-modal biometrics). In this paper, we will focus on multi-source biometrics, although most of the research results reported in this paper also apply to multi-sample systems.

"Multi Modal Technology makes Biometrics work" – so states the title of a recent press release from Aurora Defense [7]. Many other biometric vendors jump on the same bandwagon: multibiometrics is definitely a hot technology. And indeed, multibiometric systems promise significant improvements over single biometric systems, for example higher accuracy and increased resistance to spoofing. They also claim to be more universal by enabling a user who does not possess a particular biometric identifier to still enroll and authenticate using other traits, thus eliminating enrollment problems.

But can multibiometrics live up to the hype? At a first glance, incorporating multiple biometrics into one system appears to be a very intuitive and obvious concept. But as described in the next chapter, there are very different ways to actually combine multiple sources of information to make a final authentication decision. Information fusion strategies range from simple boolean conjunction to sophisticated statistical modeling.

Without going into the mathematical details, this paper reports on selected recent approaches. Our goal is to analyze how well multibiometric systems are able to keep up with the vast promises made by their advocates.

2. SYSTEM ARCHITECTURES COMPARED

As suggested in the literature (e.g. [3] or [4]), multibiometric systems are categorized into three system architectures according to the strategies used for information fusion:

- Fusion at the **Feature Extraction Level**
- Fusion at the **Matching Score Level**
- Fusion at the **Decision Level**

That is, we classify the systems depending on how early in the authentication process the information from the different sensors is combined. Biometric authentication is a chain process, as depicted in Figure 1 (see [2] for a more detailed explanation):


Figure 1: Authentication Process Flow

Fusion at the feature extraction level stands for immediate data integration at the beginning of the processing chain, while fusion at the decision level represents late integration at the end of the process.

The following sections describe each of these architectures in detail and report on related research activities.

2.1 Fusion at the Feature Extraction Level

In this architecture, the information extracted from the different sensors is encoded into a joint feature vector, which is then compared to an enrollment template (which itself is a joint feature vector stored in a database) and assigned a matching score as in a single biometric system (see Figure 2).
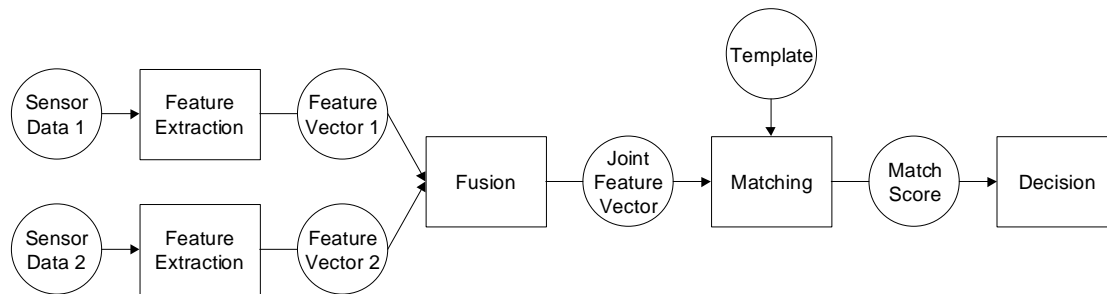

Figure 2: Fusion at the Feature Extraction Level

Even an extensive literature search did not reveal any significant recent research on this fusion strategy. This suggests that fusion at the feature extraction level is much less preferable than the other two strategies. I can identify two main problems with this approach:

- the feature vectors to be joined might be incompatible (e.g. due to numerical problems), or some of them might even be unavailable (e.g. in cases where the user does not possess all biometric identifiers). While the first issue might be resolved by careful system design, leading to a very tightly coupled system, the second one will cause the enrollment problems we already know from single biometric systems.

- score generation is problematic: even in a single biometric system, it is difficult to find a good classifier, i.e. to generate a representative score based on the matching of feature vector and enrollment template. But for the high-dimensional joint feature vectors in a multibiometric system, it is even more complicated. As pointed out in [6], the relationship between the different components of the joint feature vector may not be linear.

2.2 Fusion at the Matching Score Level

In a multibiometric system built on this architecture, feature vectors are created independently for each sensor and then compared to the enrollment templates, which are stored separately for each biometric trait. Based on the proximity of feature vector and template, each subsystem now computes its own matching score. These individual scores are finally combined into a total score, which is handed over to the decision module. The whole process is shown in Figure 3:
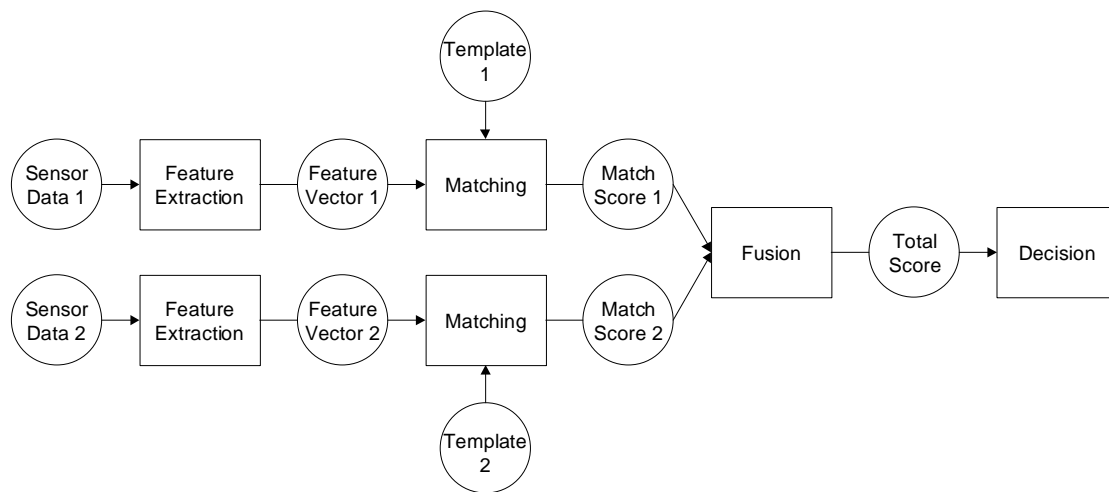

Figure 3: Fusion at the Feature Extraction Level

The process flow inside a subsystem is the same as in a single biometric system, thus allowing the use of proven algorithms for feature extraction and matching.

A very elegant example for this fusion strategy has recently (2002/2003) been presented by Ross and Jain in two research papers ([4] and [5]):

They incorporate facial scan, fingerprint verification and hand geometry scan into a common authentication system, using well-known methods for each identifier (eigenfaces for the facial scan, minutiae patterns for the fingerprint system and commonly used hand geometry features).

Matching scores for the three modalities are then normalized and combined using one of the following strategies:

- The **Sum Rule** is to take the weighted average of the scores.

- The **Decision Tree** strategy uses a sequence of threshold comparisons on the different scores to make an authentication decision. According to the authors, the thresholds were computed using the tree based machine learning software C5.0 to maximize information gain for each comparison.

- The **Linear Discriminant Analysis** transforms the 3-dimensional score vectors into a new subspace, in which the separation between the classes of genuine user scores and impostor scores is maximized. The optimal parameters for this transformation are calculated in advance based on a training data set. The output score is defined as the minimum distance to the centroids of the two classes, using a special metric, the Mahalanobis distance.

Based on experimental results, the authors make the observation that the sum rule achieves the best performance. Most importantly, they further extend the sum rule using a really new approach: they suggest applying user-specific weights to the individual traits to be combined as well as using user-specific threshold levels for making the final authentication decision.

The authors also present the corresponding learning rules: initially, equal weights are assigned to each biometric trait, which are then varied after each use to minimize the sum of the false accept and false reject error rates. For the thresholds, each user's cumulative histogram of impostor scores for the different biometric identifiers is used. Unfortunately, the authors do not give further details, neither do they present alternative learning rules, which might perform even better, e.g. neural networks or other machine learning approaches.

Nevertheless, this strategy of user-specific weights is certainly the best solution I have seen so far to deal with non-universal biometric traits and enrollment problems. If a user does not posses a certain biometric identifier or shows only weak characteristics, the corresponding weight can be adjusted to a small value.

The final question to be answered is whether this approach really leads to a higher accuracy. And indeed, the experimental data that the authors provide suggests pretty good performance for the combination of all three biometric identifiers. However, it is not significantly better than the best fingerprint systems tested in [8]. This might be due to the fact that the individual subsystems used in this experimental system are rather weak, especially their hand geometry verifier. We can therefore hope to achieve even better performance when combining top-of-the-line verifiers for each biometric trait.

2.3 Fusion at the Decision Level

In this fusion strategy, a separate authentication decision is made for each biometric trait. These decisions are then combined into a final vote, as shown in Figure 4:
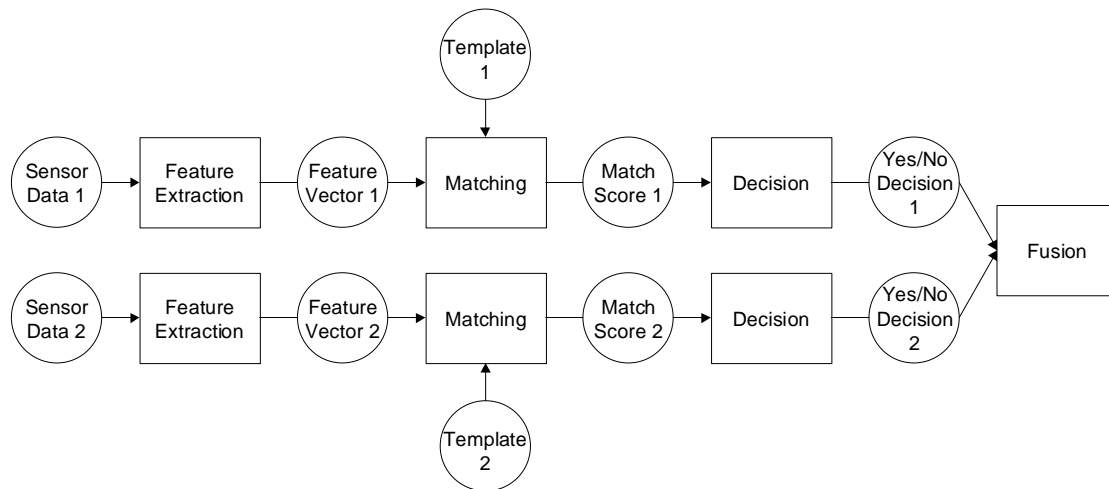

Figure 4: Fusion at the Decision Level

Fusion at the decision level is a rather loosely coupled system architecture, with each sub-system performing like a single biometric system. This architecture has therefore become increasingly popular with biometric vendors, often advertised under the term "layered biometrics". The emergence of biometric standards like BioAPI ([9]) has further supported this concept.

Many different strategies are available to combine the distinct decisions into a final authentication decision. They range from majority votes to sophisticated statistical methods (e.g. as described in [6]). In practice, however, developers seem to prefer the easiest method: boolean conjunctions.

The renowned BioNetrix Authentication Suite, for example, offers the following combination strategies (among others; see [10] for the full list):

- the **AND rule** requires a positive decision from all verification modules. While this will certainly lead to low false authentication rates, it will also result in high false rejection rates.

- the **OR rule** attempts to authenticate the user using one biometric trait. If this fails, he is offered another attempt with another verification module. This policy is trading a low false rejection rate for a high false authentication rate.

- a very interesting rule is the **RANDOM rule**, where a biometric trait is randomly chosen. Although this is a very simplistic idea, it can definitely make it harder for intruders to spoof the system. But it comes without the inconvenience of a multi-level data acquisition for each authentication attempt.

Fusion at the decision level occurs at a very late stage of the authentication process. We can therefore assume that it does not show the same potential to improve the overall system performance as fusion at the matching score level. Only under very specific conditions, accuracy improvements can be guaranteed [3]. As Daugman shows, if these conditions are violated by using biometric tests which differ significantly in their performance, their combination at the decision level can even lead to serious performance degradation.

3. EFFECTS OF MULTIBIOMETRICS ON THE USER

So far, we have only dealt with the internal architecture of a multibiometric system. The effects of multibiometrics on the user are not discussed in any of the references I have found. In [3], Hong et al. even make the amusing statement: "Finally, we assume that offering multiple biometric identifiers presents a negligible inconvenience to the user".

Is this assumption justified? Are multibiometrics only a "negligible inconvenience"? First and foremost, I can see major privacy issues tied to multibiometrics. In a multibiometric system, the user has to reveal a whole spectrum of biometric identifiers, with all of them being stored in the template database after the initial enrollment. The user profiles stored in such a database are therefore significantly more comprehensive than in a single biometric system. Hence, it becomes a very attractive target for identity thieves. Biometrics vendors repeatedly claim that the original data can not be restored from the enrollment templates. However, we have no way of verifying, since the feature extraction algorithms are always proprietary and never made available to the public. And in fact, there are commercially available systems for which the contrary has been shown (e.g. in [12]).

In traditional password-based authentication, a user can simply choose a new password once the old one has been compromised. It is a major problem of biometrics that most biometric identifiers cannot be changed. If, for example a fingerprint is compromised, you cannot just get a new one. In this case, you would need to switch to another biometric identifier, e.g. use another finger. But in a multibiometric system, this one might be compromised as well, which makes the problem even more severe.

Another problem I can clearly identify is the inconvenience of a multi-level data acquisition process to the user. The different biometric identifiers can either be obtained sequentially or simultaneously, but both ways have their disadvantages: if they are acquired one after another, it will take considerably more time for users to authenticate and thus reduce productivity. On the other hand, if you have already used biometric systems, you know that these rely on good data quality and are therefore sensitive to factors like positioning, clarity of voice, etc. It might be a challenge to provide good samples for multiple biometric identifiers at the same time. Imagine how funny it would look if a user tried to position his thumb on the scanner, while at the same time rotating his head to pass the face recognition.

Of course, not all multibiometric systems will be equally affected by the problems mentioned above. A good system design as well as a careful choice of the biometric traits to be used can certainly alleviate some concerns. And we should not forget about the obvious advantages which multibiometric systems may offer to the user, such as lower failure-to-enroll rates and higher accuracy of authentication. It is still too early to predict whether these will be sufficient to make users accept the inconveniences. But in any case, all possible effects on the user should be discussed openly. At the moment, this is still not happening, with adverse affects being left out as pointed out at the beginning of this chapter.

4. CONCLUSION

"Multi Modal Technology makes Biometrics work" – this was the advertising slogan that we have started with. We have discussed several different approaches to multibiometric systems. And indeed, we have encountered interesting attempts to alleviate some of the problems from which conventional biometric systems still suffer. The most promising recent research is certainly the information fusion at the matching score level involving user-specific weights and threshold levels, as suggested by Ross and Jain. This approach might have the potential to finally get rid of the nasty enrollment problems and at the same time improve accuracy of authentication.

Furthermore, it is obvious that the simultaneous acquisition of multiple biometric identifiers makes it a lot harder for an impostor to spoof the system by presenting artificially created samples.

However, we do not get those benefits for free: multibiometric systems are less cost-effective, and they have significant effects on their users. Some of these could in fact lead to reduced user acceptance: especially the privacy issues and the inconvenience of multi-level data acquisition might cause acceptance problems.

Many of the promising architectures for multibiometric systems are still at an experimental stage. Currently available multibiometrics are mainly layered, featuring only loose coupling between the different subsystems, sometimes even with different user interfaces.

It is now up to the developers and vendors to present truly integrated solutions with higher accuracy and at the same time improved ease of use, despite multiple biometric identifiers being acquired.

REFERENCES

General Biometrics References

[1]   Bubeck, U. M. and Sanchez, D. "Biometric Authentication: Technology and Evaluation", 2003, http://www.ub-net.de/informatik/pub/biosurvey/biosurvey.pdf

[2]   Nanavati, Samir et al. "Biometrics: Identity Verification in a Networked World". Wiley Computer Publishing, New York, 2002

Multibiometrics References

[3]   Hong, L. et al. "Can Multibiometrics Improve Performance?" Proceedings AutoID 1999, URL: http://web.cse.msu.edu/TR/MSUCPS:TR99-39

[4]   Ross, A. and Jain, A. K. "Information Fusion in Biometrics". to appear in Pattern Recognition Letters, 2003, URL: http://biometrics.cse.msu.edu/RossFusion_PRL03.pdf

[5]   Jain, A. K. and Ross, A. "Learning User-specific Parameters in a Multibiometric System". Proceedings International Conference on Image Processing (ICIP), 2002, URL: http://biometrics.cse.msu.edu/JainRossICIP2002.pdf

[6]   Prabhakar, S. and Jain, A. K. "Decision-level Fusion in Biometric Verification". Pattern Recognition v35 n4, 2002, URL: http://www.cse.msu.edu/cgi-user/web/tech/document?NUM=00-24

[7]   „Multi Modal Technology makes Biometrics work". PRWeb Press Release, Aurora Defense LLC, 2002, URL: http://www.prweb.com/releases/2002/2/prweb33800.php

[8]   Mansfield, T. et al. "Biometric Product Testing Final Report". UK Biometrics Working Group, 2001, http://www.cesg.gov.uk/technology/biometrics/media/Biometric Test Report pt1.pdf

[9]   Tilton, Catherine J. "An Emerging Biometric API Industry Standard". IEEE Computer v33 n2, 2000

[10]   Speir, Michelle. "BioNetrix delivers layered biometrics suite". Federal Computer Week, 2000, URL: http://www.fcw.com/fcw/articles/2000/0605/web-biobrf2-06-05-00.asp

[11]   Daugman, John. "Combining Multiple Biometrics". Cambridge University, URL: http://www.cl.cam.ac.uk/users/jgd1000/combine/combine.html

[12]   Thalheim, Lisa et al. "Body Check: Biometric Access Protection Devices and their Programs Put to the Test". c't 11/2002, URL: http://heise.de/ct/english/02/11/114/