Overview of Issues on Standardisation of Multi-Modal Biometric Combiners

Nigel Sedgwick, Cambridge Algorithmica Limited, 16th October 2003

Some applications of biometrics require a level of technical performance that is difficult to obtain with a single biometric device. Such applications include prevention of multiple applications for national identity cards and security checks for air travel. In addition, provision is needed for people who are unable to give a reliable biometric sample for some biometrics.

Use of multi-model biometric measurements, from substantially independent biometric devices, gives improved technical performance. This includes an improved level of performance where not all biometric measurements are available, such that decisions can be made from any number of biometric devices within an overall policy on accept/reject thresholds.

Currently, there are concerns that multi-modal biometrics cannot be combined reliably to give a guaranteed improvement in performance. This is, in fact, incorrect from a fundamental theoretical viewpoint. Though there are practical difficulties, these should not be viewed as overwhelming. Any number of suitably characterised biometric devices can have their decision scores combined in such a way that the multi-modal combination is guaranteed (on average) to be no worse than the best of the individual biometric devices. In practice, there will always be an improvement from multi-modal combination.

However, the mechanism (for this sort of good combination of scores within a multi-modal biometric system) must follow certain guidelines. Firstly, each biometric device must produce a score, rather than a hard accept/reject decision, and make it available to the multi-modal combiner. Secondly, in advance of operational use, each biometric device must make available to the multi-modal combiner, its technical performance in the appropriate form (and with sufficient accuracy of characterisation).

These two aspects, of data transfer from individual biometric devices to a multi-modal combining system, each requires standardisation. Furthermore, this impinges on other aspects of standardisation, such as the BioAPI and CBEFF. There are also implications for SC37/SG5 on Biometric Testing and Reporting.

It is recommended that standardisation of multi-modal biometric combination should be treated as an additional standard with one or more additional levels of application (or accreditation). Thus, a manufacturer of an individual biometric device could make no claim of suitability for use within a multi-modal system, or could make claims of its compliance with the multi-modal standard at specified levels.

The technical approach for multi-modal combination recommended here is thought to be optimal, in terms of statistical theory, particularly Bayes theorem. It is as follows, in summary.

Scores from individual biometric devices should be expressed as likelihood ratios (of impostor over genuine). This can be done within the individual biometric device or, preferably, within the multi-modal combiner. The advantage of doing this within the multi-modal combiner is that there is the opportunity for replacement of the characterisation data of the individual biometric devices with characterisation data derived independently, rather than only being able to rely on characterisation data provided by the equipment manufacturer.

Optimal multi-modal combination is done just by multiplying the likelihood ratios from the contributing individual biometric devices. This gives a multi-modal likelihood ratio, which can itself be further combined by multiplication with other likelihood ratios. [Note. There are practical benefits from using the logarithm of likelihood ratios, and combining them by addition.]

The likelihood ratio (from one or any number of biometric devices) can be interpreted as the Biometric Gain against Impostors (BGI). The BGI tells us how many times more (or less) certain we are that the applicant is an impostor, than we were before the measurement. The

BGI has the particular advantage that it requires no a priori knowledge of the probability that applicants are genuine or impostor.

[Note. The likelihood ratio is a very close approximation of the BGI, for circumstances where the a priori probability of an impostor is very low (as is usually the case). In this case, the BGI is effectively independent of the a priori probabilities. Where the a priori probability of being an impostor is not low (and is known) there is a simple and explicit equation for calculating it from the a priori probabilities and the likelihood ratio.]

There is a further issue that, for some individual biometric devices, it is desirable to provide the device characterisation information for each individual enrolled genuine user, rather than as a single characterisation for all enrolled users. This has further implications for standardisation of interfaces and data formats.

The characterisation data that is required for each individual biometric device are the two Probability Density Functions (PDFs), of scores from genuine users and of scores from impostors. [Note. These PDFs can be derived from the same experimental data that is used to derive ROC measurements, which is nearly always readily available.] Where characterisation is required for individual enrolled users, these 2 PDFs are required for each enrolled user. In practice, the PDFs used for calculating likelihood ratios are likely to be characterised by mathematical functions with a small number of parameters.