Signature verification revisited: promoting practical exploitation of biometric technology

by M. C. Fairhurst

Despite research over a long period, biometric approaches to authenticating personal identity have not met with the degree of success in practical applications originally predicted. This paper discusses approaches to biometric testing, focusing particularly on automatic signature verification, and addresses some of the important issues which might help to promote the introduction of practical systems in the future. Examples associated with the adoption of an explicitly flexible approach to signature verification are used to illustrate the discussion, and it is argued that there is still considerable potential for practical exploitation of this type of technology.

1 Background

In an increasingly complex and sophisticated society, more and more demands are being placed on individuals to be able, when challenged, to produce proof of identity. In fact, so widespread, pervasive and ingrained is this concept that in all but the most extreme or unusual circumstances most people do not even recognise a situation as a 'challenge', and would certainly not consider the procedure at all out of the ordinary. Hence we find it perfectly acceptable to tap in a digit-based PIN (personal identification number) identifier to obtain money from a bank cash dispenser, to sign a sales voucher for scrutiny by a shop assistant, or to have our passport photograph checked at an airport.

The purpose of such procedures is to provide corroborating evidence to authenticate our claimed identity — to confirm we *are* who we *say* we are — and it is clear, even from these examples, that the evidence we can supply can take several different forms. The most common authenticating strategy is to identify a physical entity which itself is of known provenance (e.g. the possession of an approved swipecard for access control), or to check guaranteed information known on a restricted basis (PIN number, password, etc.). A third common option is to validate identity by means of a specific individual activity pattern — the basis of the legal handwritten signature or of voiceprints — and, finally, it is possible to look for recognisable and unique physical features possessed by an individual (fingerprints are the most familiar, but iris scanning or hand geometry are other possibilities)[1–5].

Each possible approach has its own positive and negative characteristics which are suited to a greater or lesser extent to different applications and situations, and each offers reliability of a greater or lesser degree. All rely on increasingly sophisticated technology for their introduction and regular use. Most of the techniques falling into the first two categories of the four noted above are already in widespread use but are also unfortunately susceptible to compromise on a large scale (e.g. PIN numbers are very frequently passed on or written down for ease of recall, entry cards and other documents are easily copied, stolen or lost).

For this reason there has been a long-standing interest in developing methods of personal identity authentication based on strategies falling into the third and fourth categories. These techniques all rely on biometric measurements, defined by the Association for Biometrics[6] as:

> 'the automated measuring of one or more specific attributes or features of a person, with the aim of being able to distinguish that person from all others'.

It is clear that we can subdivide the biometric options further, broadly based on the two categories noted above, into *physiological* biometrics (iris characteristics, fingerprints, etc.) and *behavioural* biometrics (signature checking, keystroke dynamics, etc.).

Although many early devices purporting accurately to measure biometric characteristics often did not live up to expectations, more recent developments have made viable, in principle, the introduction of technology which, under appropriate circumstances, can lead to commercial exploitation of biometric testing. The two principal issues relating to the introduction of a specific biometric authentication procedure concern the reliability and robustness of the procedure (essentially a technological issue) and the extent to which the intended user community finds the technique and its implications acceptable (largely a social issue).

Both issues can generate heated debate but, whereas the first is very much open to discussion, in the second area there would be a general consensus that, although alternatives can of course be considered, the use of the personal handwritten signature has the significant advantages that:

(*a*) the signature is the most natural and generally established of all the ways in which we seek to confirm our identity. This is confirmed in a number of studies, including one carried out by the author[7].

(*b*) the use of signature verification will minimise the disruption to accepted practices with respect to transactions where personal identity has to be authenticated

(*c*) measurement of signature characteristics is non-invasive (compare this with other potential techniques such as iris scanning) and has no negative or undesirable health connotations (as might be the case with, say, fingerprint checking, which is often considered to raise civil liberties issues and which, in use, involves direct physical contact with a possibly contaminated surface). This is potentially a most important practical issue, the significance of which should not be underestimated.

This paper will examine some of the important issues underlying the exploitation of technological approaches to the processing of handwritten signatures as a means of verifying personal identity. Although reviewing some of the underlying technologies which may be used to address the problems of signature verification, the paper will particularly seek to show how the way in which a system is configured and matched to a specific set of task requirements is equally influential in developing a viable and successful system.

Taken together, these aspects of the design and implementation of a signature verification processor will point the way to providing the major components and processing environment which will best promote the use of automatic signature verification in practice, and it is this which is the principal focus of this paper.

2 Approaches to automatic signature verification

Approaches to the signature verification problem can be broadly categorised into static verification methods, dynamic verification methods (where both rely on underlying statistical or structural processing methodologies), or neural network techniques. The static approach, based as it is on the limited information available from the shape and structure of the signature image alone, represents a very difficult problem. Techniques can be found which are based on the extraction of shape descriptors or on stroke models based on psychophysical observations of signature execution[8]. Recent work has adopted a technique based on relational similarity measures between reference profiles and pairs of text primitives[9].

Many dynamic approaches are based on dynamic programming techniques, and appropriate signals for measurements are well known[10]. Again, a variety of subsidiary paradigms have been invoked. Some are based directly on models of the handwriting generation process[11], or on signature analysis using perceptually important points[12], while others seek a multilevel approach to focus particularly on the elimination of forgeries[13]. Recent work[14] has used spectral stroke analysis in conjunction with a structured knowledge base. Some systems, such as the KAPPA approach developed by the author[15], are able to integrate efficiently dynamic and static information.

Recently, neural network approaches have been generating significant interest, and have been used in segmentation and in tackling static and dynamic problems[16–18]. Despite the many advantages of this paradigm, a significant disadvantage is generally the requirement for large training sets of samples, which precludes application in most practical problems.

Although it is therefore possible to address the problem of automatic signature verification in a variety of ways, such as by treating the signature as a time-varying signal which can be matched to a reference template, it is more common and generally more effective to approach the problem by extracting specific features from the signature image and/or execution pattern and to use these as the basis for comparison against a stored model which can reflect the statistical variations in individual signature specimens.

Since automatic signature verification has a long history it is not surprising that a wide range of features have been proposed by means of which to characterise an individual signature. For illustrative purposes, some typical examples might include measurements based on the following:

Static features:
Number of component strokes
Ratio of long to short strokes
Curvature measurements
Segment lengths
etc.

Dynamic features:
Timing measurements
Stroke order
Pen velocity profiles
Pen acceleration profiles
Pen up/pen down patterns
etc.

Likewise, signature sample acquisition is possible using a variety of different physical devices. For off-line capture any standard commercial device (e.g. document scanner, camera) may be used, since in these conditions the signature sample is represented solely as a two-dimensional image. For on-line capture, it is customary to use a digitising tablet with associated pen. Such a device allows not only the reconstruction of the signature (through the capture of the co-ordinate stream traced out by the pen) as a visual image, but also allows the extraction of the dynamic features of signature execution. Again, there are many commercial devices available. Some require the use of a wired pen (which may sometimes constrain the signing process in an undesirable way); others utilise a free pen which creates a much more familiar signing environment for the user and causes little interference in the signing process. All such devices inherently provide information about the pen trajectory across the available writing surface, and some additionally detect pen pressure and hence offer a potential source of further dynamic information. An example of a typical device is shown in Fig. 1.

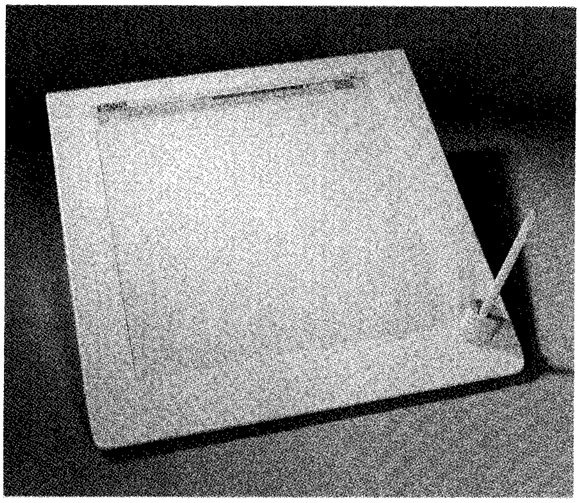

Fig. 1 A digitising tablet and pen — a typical input device for signature capture

In general, automatic signature verification requires two distinct processing stages. Initially, a potential user enrols on to the system by providing signature samples on the basis of which a class model can be constructed which (statistically) encapsulates the characteristics of an individual signer. Subsequently, a verification processing procedure is invoked to judge the likely authenticity of a presented sample with respect to its alleged class model. The process is summarised in Fig. 2.

Verification errors occurring in a signature verification system may be categorised as one of two types. On the one hand, a genuine signer may be rejected by the system as a potential forger (as could happen, for example, because of the execution of an atypical or particularly careless sample from a genuine signer), resulting in what is denoted a Type I error. On the other hand, a skillful forger might be able to generate a sample which would be accepted as genuine, resulting in a Type II error. It is seen that the verification process may be characterised as the evaluation of a discriminant function for a given test sample which can be compared against a threshold value. Whether or not the sample is considered to be acceptable or not is then determined by whether the discriminant value generated by that sample is greater than or less than the threshold chosen. The threshold setting is clearly instrumental in determining the limits of acceptability within which the signature will be considered to be genuine, and it is this which will define exactly the nature of the trade-off which can be achieved between maintaining a resistance of the system to the possibility of compromise through forgery and the very real practical problem of ensuring that the number of false rejections generated is as small as possible.

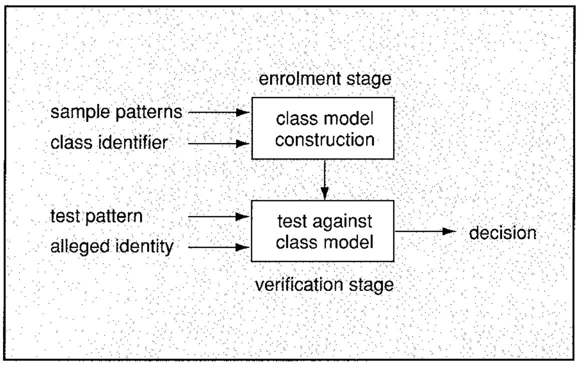

Fig. 2 Processing stages for automatic signature verification

The approach commonly taken in designing a verification system, which adopts the principle of seeking a universally applicable set of features on which to base a judgement about the acceptability of individual samples, has provided a wide variety of possible solutions to the problems inherent in the processing of the highly variable data associated with verification based on behavioural biometrics and the requirements of robustness and security (see, for example, References 19 and 20). However, it is increasingly being recognised that this principle of 'universality' may be a limiting factor in reliable system design and that a greater degree of success might be achievable if the specific requirements of a given task domain and the precise way in which the system is to be used can be taken into account.

This suggests an approach to the design of signature verification systems in which considerably greater flexibility is incorporated. There are a variety of operational parameters which may be important in characterising and evaluating a system, including performance indicators such as the error rate achieved (more precisely, the balance between Type I and Type II errors), processing speed attainable, number of training samples required to construct a signature model, the degree of security afforded to the signature data, cost of implementation, robustness in operation, and so on. In many cases, performance indicators can be traded one against another, and an optimal configuration is task-dependent. Hence it appears potentially productive to adopt an approach to system design based on the development of a 'toolkit' of processing modules which, through selection and configuration in an appropriate way, can satisfy a wide range of practical requirements.

For the purposes of illustration, it is helpful to consider a practical system. An example of a system which is specified using exactly this type of flexible implementational framework is the KAPPA signature verification system[15]. This is a system which can operate in either on-line or off-line mode, can make independent checks of static and dynamic features of the signature, and can draw on a variety (unlimited in principle) of feature types to drive the verification algorithms. The system is structured in a very flexible and adaptable way, and hence can be used in a variety of application areas with differing practical requirements.

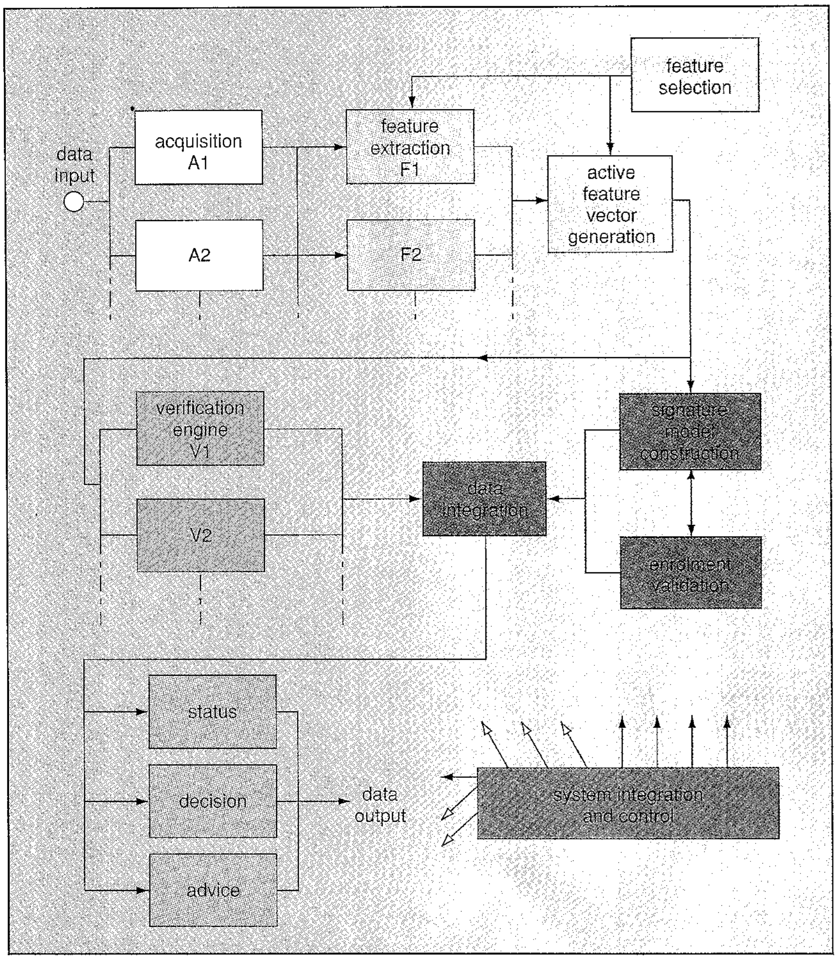

Fig. 3 Structure of the KAPPA signature verification system

KAPPA embodies a number of characteristics of practical significance. The system has a highly modular structure configured as a design toolkit, it can achieve on-line and off-line processing, it can utilise multisource feature extraction (depending on the data capture facilities adopted), it has the capacity for feature selection, offers many optimisation features such as enrolment model validation, and can provide robust and efficient operation in a variety of task domains. An overview of the broad structure of the system, emphasising its modularity, is shown in Fig. 3, and Fig. 4 shows a typical PC-based implementation.

This type of system is very useful in illustrating some fundamental principles for maximising the effectiveness with which automatic signature verification can be made to work in practice. In particular, it is interesting to consider how the tools provided can be used to optimise system performance in appropriate circumstances. Some of the important issues to be considered in this context will be examined in the following Sections from a practical point of view.

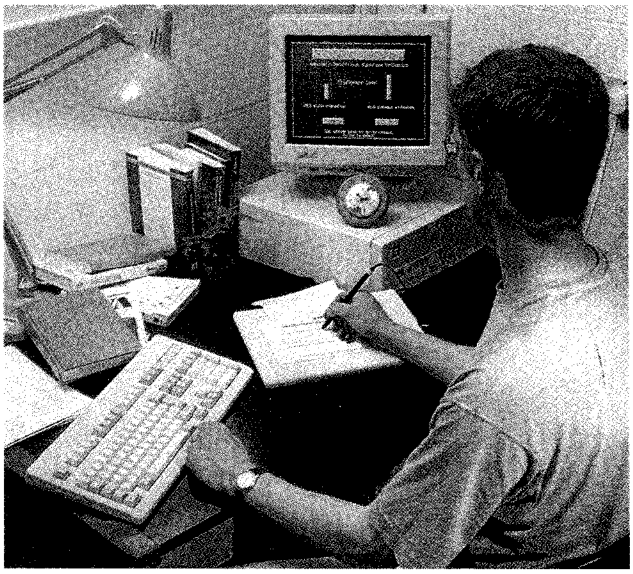

Fig. 4 A PC-based implementation of the KAPPA system

3 Feature selection for system optimisation

As noted above, conventional approaches to the problem of automatic signature verification generally adopt a procedure for constructing a universal or class-independent feature vector which then forms the basis of a verification decision using an appropriate decision function. However, the highly individual nature of the handwritten signature suggests that an approach which adopts a 'personalised' optimisation of the features used for verification, i.e. where the verification process is based on the extraction of an individually selected set of descriptive features to characterise the signature of each enrolled user of the system, might be productive.

A principal difficulty with attempting meaningful feature selection for the specific problem of automatic signature verification is the generally encountered lack of adequate statistical information to describe authentic patterns. In order to select an optimum set of features requires a mechanism for generating combinations of possible features and a metric by means of which to assess the relative merits of the features to identify the optimum set, and in this application these requirements point to the need for some form of heuristic, rather than more formal analytical, evaluation of an appropriate criterion function for use in determining a useful feature set.

An approach embodying this principle has been developed based on the progressive generation of an optimised active feature set for each individual signer determined by invoking a novel criterion function which is locally computed from an analysis of the relation between Type I and Type II errors as a function of the applied discrimination threshold. Here, error performance for a given individual signer is evaluated on the basis of a reference population of possible signers. Fig. 5 shows the principle behind the measurement of the criterion function, where individualised feature selection begins from a pool of possible features from which an algorithmic selection procedure evolves an 'active' subset of features for an individual signer until optimality is attained. Here, the criterion function is chosen as the separation between the respective Type I/Type II zero point crossings on the verification threshold axis (denoted the '0% region'), and it is this separation distance which is to be maximised.

An iterative search method can be adopted, whereby features are selectively added to and deleted from a subset of active features, depending on their relative merits with respect to the previously selected features. Specifically, a simple technique has been implemented[21] based on the evolution of the initial empty active feature vector through a process of iteratively adding n features and subsequently deleting m features ($n > m$) until the criterion function has satisfied a termination criterion. Fig. 6 shows the typical profile of the development of the active feature vector. If too few features are available there is an under-characterising of information available to the verification processor. Also a point is reached where redundancy can again diminish the power of the active feature set in discriminating between genuine and false signature samples.

The pool of features from which the active feature vector is constructed will generally be expected to be large (typically of the order of perhaps 100 or more features), and this leads to a situation where the generation of the individually optimised feature vector incurs very large computational overheads. Although a direct serial implementation of this type of approach is consequently precluded in almost all practical cases, it has been found that a very efficient mapping to a parallel hardware implementation is possible[22]. In terms of achievable performance, this type of optimisation approach can be very effective in appropriate circumstances. For example, in a small-scale trial with a group of 22 signers each donating 40 sample signatures, it was found that an error rate of 1·82% when operating with a universal feature vector of 40 elements was reduced to 0% with implementation of the individual feature optimisation procedure. An additional advantage is the significant potential reduction in the number of features which must be derived from each sample during the operational phase and which must be stored at some point in the enrolment/verification processing chain.

4 System optimisation at the operational level

A further advantage of the modular approach may be illustrated conveniently by considering a situation where operational constraints within a specific task domain can be controlled.

In the simplest case, it can be seen that there are various situations in which it is possible to disregard a 'single shot' false rejection of a valid signature by introducing the possibility of a 'retry' facility, often employed in other similar situations (for example, in PIN-based ATM access), where the signer is allowed, say, three attempts to generate an acceptable signature sample before overall rejection occurs. This is clearly a decision which can be made purely at the operational level, but one which, where appropriately introduced, can significantly reduce the problem of a higher-than-acceptable false rejection rate without necessarily compromising the security of the system to attempted forgery in a significant way.

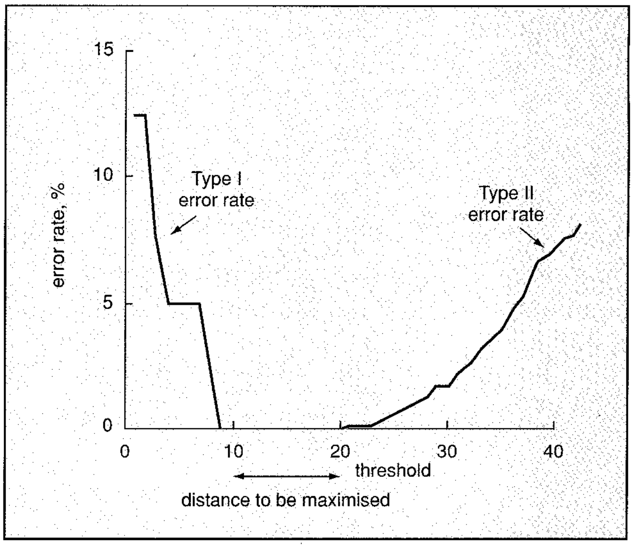

Fig. 5 The principle behind measurement of the criterion function

More generally, one of the modular components of the KAPPA system is an enrolment validation module, the purpose of which is to ensure that false rejections do not occur simply as a result of using a signature model which is inadequate or unrepresentative of the range of signature specimens which are validly generated by a particular signer. This is an important issue, mainly because the inherent nature of a practical signature verification system is that operational conditions generally preclude the availability of a large set of samples from which to construct the signature model, leading to a significant risk that the model does not adequately represent the signing profile of any given individual. By invoking an enrolment validation module it is possible to construct an initial signature model on the basis of a (fixed) small number of donated samples, but to seek additional samples with which to refine the

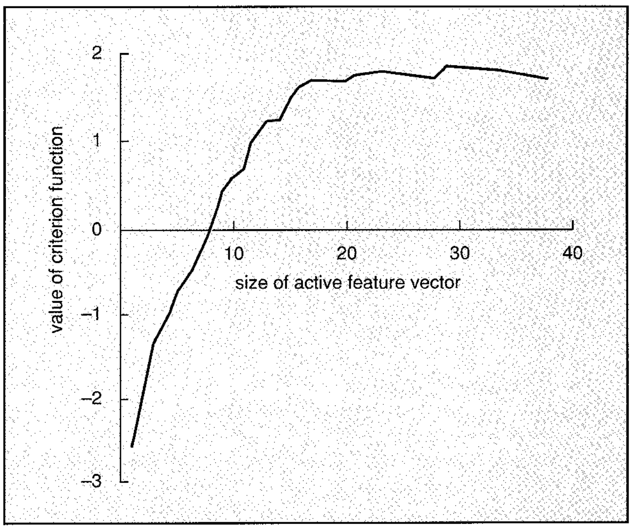

Fig. 6 Typical profile of the development of the active feature vector

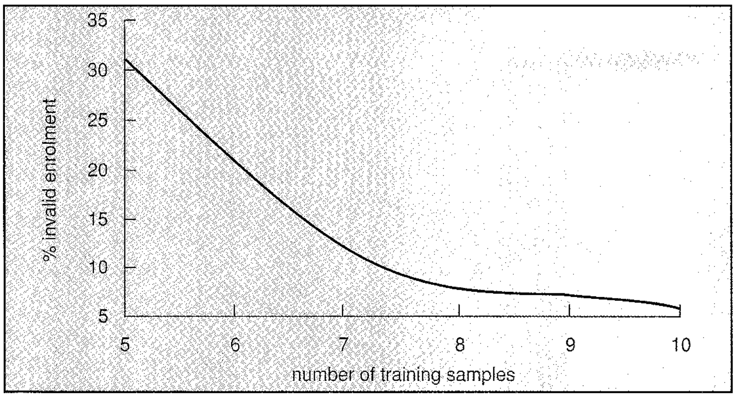

Fig. 7 Effect of number of training samples on enrolment validation in an on-line verification process

model if its initial representation is shown to be unsatisfactory. This allows a greater degree of control over the characteristics of the model ultimately adopted when the system is routinely used for verification of unknown samples.

This idea is illustrated in Fig. 7, which shows the effect of enrolment validation in an on-line verification process when up to ten donated samples are allowed for construction of an individual signature model. It is apparent that an enrolment set size based on an initial small number of donated samples (five in this case) frequently leads to an unsatisfactory model, thereby potentially seriously compromising the subsequent performance and reliability of the system as a whole. In fact, as can be seen, using this fixed and limited set of samples for enrolment leads to a situation where around 30% of enrolees would have subsequent signatures processed with an unsatisfactory reference model, thereby significantly increasing the risk of compromising the reliability of system performance. If, however, further samples (up to a maximum of ten in the case shown here) are allowed, only around 6% of enrolees would then fail to enrol satisfactorily.

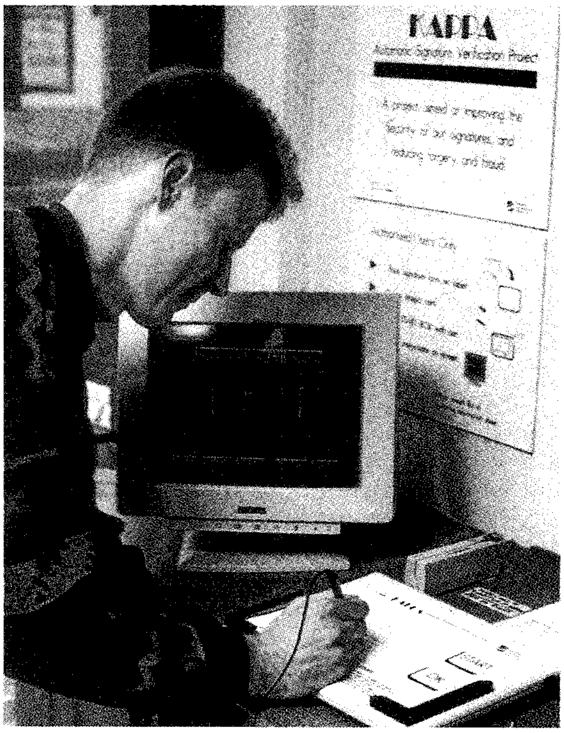

A prototype implementation of the KAPPA system in use in trials conducted at Hedge End Post Office, near Southampton. These trials allowed the collection of more than 8 000 signature samples from more than 300 signers. The trials took place in a typical retailing environment, involved a true cross-section of the general public, and allowed signatures to be collected over an extended period of several weeks. The database of signature samples collected in this way is especially valuable, since it reflects the characteristics of a population which can be considered highly representative of a typical operational environment in which a practical verification system might have to work. The photograph shows a volunteer donating samples during the trial.

More generally, the overall system error was measured using a large database (comprising more than 8 000 signature samples) generated during public trials with the system[2]. The results of this testing bear out the importance of reference model validation, for it was found that using a fixed enrolment sequence of 5 samples per enrolee resulted in a false rejection rate as high as around 20%, whereas allowing up to 10 samples per enrolee reduced this to around 6·9%. Furthermore, if individuals who failed to provide a satisfactory enrolment were then excluded from the system, the error rate fell further to around 1·8%. Finally, allowing a *retry on rejection* facility (up to 3 attempts) allowed the system to operate at a false rejection rate of significantly less than 1%. It is clear, therefore, that options such as this, which can be conveniently and naturally introduced within a modular framework, can be extremely valuable in realising levels of performance which are likely to encourage the introduction of this type of technology into practical situations.

5 Static enrolment validation

The previous Section described how the utilisation of an enrolment validation facility can be optimised through the introduction of simple operational procedures which can be controlled within a known operating environment. This focused specifically on a situation where on-line processing enabled dynamic information about signature execution to be considered. There are many practical situations, however, in which only static information extraction (i.e. off-line processing) is possible, and it is clear that the incorporation of a similar validation procedure could potentially be very effective in such situations.

In these circumstances there are two fundamental difficulties to be addressed. The first is the obvious fact that static features, by their very nature, are likely to provide a much less rich source of verifying data on which the processing module can operate. The second is that in most situations where static processing is necessary, the number of enrolment samples which can be utilised is generally strictly limited by the nature of the specific task.

However, it is possible to show that the principle of enrolment validation is still applicable even in these situations. For example, experiments have been carried out to develop an algorithm which can evaluate a set of signature samples, rejecting samples which degrade the current reference model. Reference models were constructed from a pool of 15 candidate samples per signer for a group of approximately 250 signers on the basis of selecting the best *n* out of *m* samples. In order to illustrate the global effect of this procedure, performance was measured on the basis of the so-called 'equal error rate', found by plotting the respective changes in Type I and Type II individual error rates as a function of acceptance threshold, and recording the point at which the two error curves intersect (i.e. the point at which the Type I and Type II error rates are equal). Some test results are illustrated in Figs. 8 and 9, which compare the performances attainable with and without this type of enrolment validation.

A number of interesting points become apparent from these results. It is immediately clear, for example, that the verification performance achievable is limited when only static features can be extracted from the signature samples, but that the introduction of an enrolment validation procedure is effective in considerably improving the error rates generated. Thus, in static processing, the introduction of such a mechanism may often be particularly important. Figs. 8 and 9 also illustrate clearly the difficulty of operating with a small number of enrolment samples, and it is seen that typically at least 8–10 valid samples per enrolee are required in order to achieve performance stability. This type of analysis clearly emphasises the desirability of adopting some form of enrolment validation, particularly in situations where only static feature extraction is possible, or in situations where small numbers of signature samples are available from which to construct a reference model.

6 Conclusions and a look to the future

Although research in automatic signature verification has a long history, attempts to design a system sufficiently robust to allow the widespread introduction of such technology in practice have not generally met with the degree of success originally anticipated. It has been claimed by potential users that the performance levels achieved to date have been the principal reason for this, yet the nature of the data generated in this type of application introduces a limiting factor which considerably reduces the scope of what can be done to achieve the levels of performance often quoted as being required. The same has been true, though sometimes for reasons of public acceptability rather than achievable performance, of other alternative forms of biometric measurement.

Starting from the widely-held viewpoint that the handwritten signature is the most traditionally recognised, natural and acceptable form of biometric, this paper has sought to address a number of key issues which might be particularly influential in promoting the exploitation of signature verification technology in a practical domain. It has been shown, for example, that a universal solution to the signature verification problem may be to all intents and purposes impractical, while better matching to specific task domains might be more productive. To this end a system architecture has been developed which offers a toolkit for a system designer to facilitate this type of task-oriented optimisation.

Likewise, it is clear that — apart from focusing on the inherent merits of individual signature processing algorithms themselves — an examination of other significant operational factors can lead to approaches for improving attainable levels of performance in practice, and that a clear understanding of the limitations and constraints of a given operational environment can be instrumental in maximising the benefits afforded by technological solutions. It is apparent, however, that for many applications, particularly those which seek large-

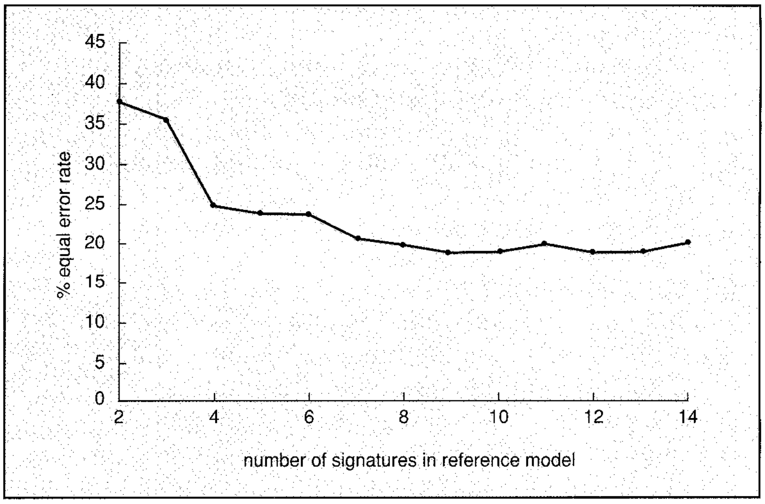

Fig. 8 Equal error rate for non-validating reference model

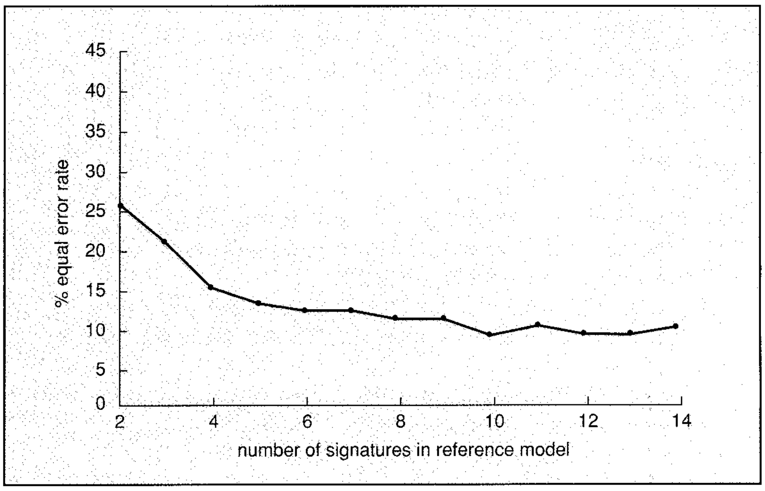

Fig. 9 Equal error rate for validating reference model

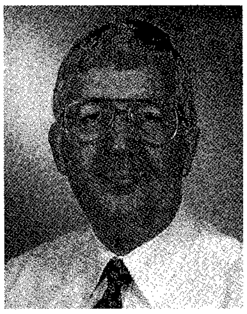

Professor Michael Fairhurst has been on the academic staff of the Electronic Engineering Laboratory at the University of Kent since 1972. He has been actively involved in various aspects of research in image analysis and computer vision, with a particular interest in computational architectures for image analysis and the implementation of high-performance classification algorithms. Application areas of principal concern include handwritten text reading and document processing, security and biometrics, and medical image analysis. Professor Fairhurst is a current member and past Chairman of the IEE Professional Group E4 on Image Processing and Vision, and has in the past been a member of the Professional Group Committee for Biomedical Engineering. He has been Chairman of several of the IEE series of international conferences on Image Processing and Applications.

Address: Electronic Engineering Laboratory, University of Kent
at Canterbury, Canterbury, Kent CT2 7NT, UK.
Email: M.C.Fairhurst@ukc.ac.uk

scale introduction of a new technology, automatic signature verification solutions may be resisted and it is unlikely that signature verification, in common with most other currently available biometrics, will be able to offer the absolute levels of attainable performance which have been specified to date.

It is possible, however, that these stringent requirements will be reviewed, but in the meantime there are still positive messages for those who believe that practical biometric testing still has much to offer even in the short term. The first is that the sort of modular design approach advocated here should allow the introduction of more reliable products than might otherwise be possible . The second is that there are an increasing number of smaller-scale, localised application areas where existing technology can find an effective role to play. The final point, and perhaps that which is ultimately most likely to be of real benefit, is that great potential exists for the introduction of systems which integrate multiple biometrics or which use biometric testing as a secondary or confirmatory check in a broader-based protocol for establishing or validating individual identity. Far from biometrics being the stuff of science fiction (or even the research laboratory), existing technology may yet offer opportunities for practical success here and now.

References

1 NAKAMURA, O., MATHUR, S., and MILHAMI, T.: 'Identification of human faces based on isodensity maps', *Pattern Recognit.*, 1991, **24**, pp. 263–272

2 HRECHAK, A. K., and McHUGH, J. A.: 'Automated fingerprint recognition using structural matching', *Pattern Recognit.*, 1990, **23**, pp. 893–904

3 DAUGMAN, J.: 'High confidence recognition of persons by rapid video analysis of iris texture'. Proc. ECOS95 (European Convention on Security and Detection), *IEE Conf. Publ. No. 408*, 1995, pp. 224–251

4 ASHBOURNE, J.: 'Practical implementation of biometrics based on hand geometry'. *IEE Colloquium Digest No. 1994/100:* Image Processing for Biometric Measurement, pp. 5.1–5.6

5 GEORGE, M.: 'A new strategy for low power, high discrimination voice biometrics'. Smart Card 1995, pp. 130–134

6 Association for Biometrics (Contact: P. Hawkes, BTG plc, 101 Newington Causeway, London, SE1 6BU): Workshop on 'State of the art in signature verification', NPL, Teddington, November 1993

7 KAPPA signature verification public trials and public survey on biometrics, BTG plc, 101 Newington Causeway, London, SE1 6BU, May 1994

8 LEE, S., and PAN, J. C.: 'Offline tracing and representation of signatures', *IEEE Trans.*, 1992, **SMC-22,** pp. 755–771

9 SABOURIN, R., and PLAMONDON, R.: ' Observability and similarity in spatial relations in the structural interpretation of handwritten signature images'. Proc. 7th Scandinavian Conf. on Image Analysis, Arlborg, 1991, pp. 477–485

10 PLAMONDON, R., and PARIZEAU, M.: 'Signature verification from position, velocity and acceleration signals'. Proc. 9th Int. Conf. on Pattern Recognition, Rome, 1988

11 PLAMONDON, R.: 'A model-based dynamic signature verification system'. Conf. on Fundamentals in Handwriting Recognition, Chateau de Bonas, France, 1993, pp. 75–93

12 BRAULT, J. J.: 'Segmenting handwritten signatures at their perceptually important points', *IEEE Trans.*, 1994, **PAMI-15,** pp. 953-957

13 PLAMONDON, R., YERGEAU, P., and BRAULT, J. J.: 'A multilevel signature verification system', *in* IMPEDOVO, S., and SIMON, J. C. (Eds.): 'From pixels to features III' (Elsevier, 1992), pp. 363–370

14 IMPEDOVO, S., CASTELLANO, M., and PIRLO, G.: 'On-line signature verification through stroke analysis'. Proc. AFTEC Conf. on New Concepts in Computer Science, 1990, pp. 47–53

15 FAIRHURST, M. C.: 'Automatic signature verification: making it work'. Proc. 1st IEE European Workshop on Handwriting Analysis and Recognition, Brussels, 1994

16 LALONDE, M.: 'A neural network approach to handwritten curve partitioning'. Proc. Vision Interfaces 93, 1993, pp. 136–141

17 DROUHARD, J. P.: 'Off-line signature verification using directional PDF and neural networks'. Proc. 11th ICPR, 1992, pp. 321–325

18 BROMLEY, J., and BENTZ, J. W.: 'Signature verification using a siamese time delay neural network', *Int. J. Pattern Recognit. Artif. Intell.*, 1993, **7,** p. 669

19 PLAMONDON, R., and LORETTE, G.: 'Automatic signature verification and writer identification — the state of the art', *Pattern Recognit.*, 1989, **22,** pp. 107–131

20 LECLERC, F., and PLAMONDON, R.: 'Automatic signature verification — the state of the art 1989–1993', *Int. J. Pattern Recognit. Artif. Intell.*, 1994, **8,** pp. 643–659

21 BRITTAN, P., and FAIRHURST, M. C.: 'An approach to handwritten signature verification using a high performance parallel architecture', *in* IMPEDOVO, S., and SIMON, J. C. (Eds.): 'From pixels to features III' (Elsevier, 1992), pp. 385–390

22 FAIRHURST, M. C., BRITTAN, P., and COWLEY, K. D.: 'Parallel realisation of feature selection for a high performance signature verification system'. Proc. PACTA, Barcelona, 1992, pp. 974–982

First received 14th February and in revised form 25th July 1997