Transparent Facial Recognition for Mobile Devices

N.L. Clarke[1,2], S. Karatzouni[1], S.M. Furnell[1,2]

[1]Centre for Information Security & Network Research,
School of Computing, Communications & Electronics, University of Plymouth, Drake Circus, Plymouth, PL4 8AA

[2]School of Computer and Information Science, Edith Cowan University, Perth, Western Australia

cisnr@plymouth.ac.uk

Abstract

The growing popularity and increasing functionality of mobile devices has resulted in them becoming increasingly desirable targets for theft and unauthorised use. Key to ensuring effective protection is providing an effective and usable form of user authentication. Traditional approaches, such as the PIN, have been shown to be ineffective and inconvenient. Indeed, the intrusive nature of this form of authentication is itself a barrier to acceptance by users. This paper discusses the application of alternative authentication technologies, such as biometrics in a transparent and continuous fashion – removing the inconvenience of intrusive authentication and extending the identity verification beyond point-of-entry. Unfortunately, for this type of authentication to operate, the individual biometric techniques need to be able to operate in a transparent fashion, where the system is unable to control key factors that are required for successful operation. For instance, facial recognition performs very well when facial orientation, illumination, distance and camera quality can all be controlled. Under a transparent mode of operation, these factors cannot be so closely controlled. The paper presents a study aimed at validating the feasibility of improving facial recognition algorithms for use transparently by the means of a composite facial template. The study specifically addresses the performance issues of utilising facial images with varying facial orientations, and has shown that the usability of facial recognition algorithms can be significantly improved through the use of this composite template at a minimal expense to the level of security being provided.

Keywords: Biometrics, Facial Recognition, FERET, Transparent Authentication, Mobile Devices

1. Introduction

The mobile networking landscape has changed significantly over the last decade with a transition from large form factor telephony devices to small multi-purpose multimedia communications devices. This technology has also experienced worldwide adoption, with 2.4 billion cellular subscribers currently spread across over 200 countries (GSM World, 2006). The recent introduction of Third Generation (3G) technologies has provided the underlying mechanism for a wide variety of innovative data orientated services, with approximately one million users every day adopting these new features (Best, 2006).

By providing functionality that extends beyond telephony, the mobile device has evolved from being a cumbersome telephone to become a necessity people utilise every day, for a variety of applications. This level of functionality can be seen to be

significantly expanding, with devices today having similar processing and memory capabilities to PCs of a few years ago. Indeed, their combination of portability and capability means that handsets such as Smartphones and PDAs are likely to have an increasingly significant role as mobile computing and network access devices.

This transition introduces serious security considerations for mobile users. With the ability to access and store a wide variety of more sensitive information (such as extensive contact lists, diaries, mobile banking and location based services), the need to ensure this information is not misused or abused is imperative. Whereas the theft or loss of a device might previously have been the principal risk associated with mobile devices, unauthorised access to a device that utilises these information services will potentially result in the disclosure of a greater amount of personal information, endangering a wider variety of aspects in the user's life (which could range from personal identity theft to serious corporate loss and increasingly liability).

However, the most popular access security currently takes the form of the password or PIN; secret knowledge approaches that relies heavily upon the user to ensure continued validity. For example, the user should not use the default factory settings, tell other people, or write it down. However, the poor use of passwords and PINs has been widely documented, with a recent study showing that 34% mobile phone users not use a PIN, 45% have never changed their PIN and 26% have shared their PIN with other people (Clarke & Furnell, 2005).

There are three general categories of user authentication: something you know (e.g. passwords and PINs); something you have (e.g. tokens); and something you are (e.g. biometrics) (Smith, 2002). The aforementioned secret knowledge approaches have already been shown to be inadequate for the future needs of mobile users. The use of token-based technology to improve user authentication cannot be completely ruled out, with technologies such as Bluetooth enabling the capability for day-to-day devices such as watches and jewellery to be used as potential tokens. However, to date, token-based technology has not provided any real level of security for mobile devices, with the SIM card (itself a token) always being left in situ within the handset. Finally, the ability to authenticate users based upon unique characteristics of the person is an interesting approach as it relies on the technology not the person for reliable security.

This paper begins by introducing the concept of user authentication for mobile devices using biometrics. However, key to this concept are a number of factors designed to ensure security is increased beyond point-of-entry, with minimal inconvenience to the user. Section 2 discusses this need for transparent and continuous authentication and the different types of biometric technique that would be appropriate within a mobile device context. One such technique that of Facial Recognition is further explored in section 3, with an experiment into its application on a mobile phone being discussed in sections 4 and 5. The paper concludes by discussing the experimental findings and suggesting further areas of research.

2. Biometric Authentication for Mobile Devices

The use of biometrics, or specifically distinguishing human characteristics, has existed for hundreds of years in one form or another, whether it is a physical description of a person or perhaps more recently a photograph. Biometrics can be

divided into two categories based upon the underlying characteristic they are using: namely physiological and behavioural (Ashbourn 2000). Physiological biometrics are those using characteristics based upon a physical aspect of the body, such as a fingerprint, face, iris, or retina. Behavioural biometrics utilise the unique way in which humans behaviour to characterise and authenticate us, such as the way in which we speak, type and sign our name.

People are often the key factor and inhibitor in many security controls, where the successful interaction of the user is required in order for the control to operate effectively. As such this research project sought to remove as much explicit security interaction from the user as possible but also achieving the following objectives:

- to increase the authentication security beyond secret-knowledge based approaches;
- to provide transparent authentication of the user (within limits) to remove the inconvenience factor from authentication;
- to provide continuous or periodic authentication of the user, so that confidence in the identity of the user can be maintained during usage of the device rather than simply at switch on;
- to provide an architecture that would function (to one extent or another) across the complete range of mobile devices, taking into account the differing hardware configurations, processing capabilities, and varying levels of network connectivity.

This can be achieved through utilising a combination of secret knowledge and biometric-based techniques within an appropriately flexible framework. The framework – called NICA (Non-Intrusive Continuous Authentication)[1] operates by initially providing a baseline level of security, using secret knowledge approaches, which progressively increases as the user interacts with their device and biometric samples are captured. Although user authentication will begin rather intrusively (e.g. when the device is switched on for the first time), with the user having to re-authenticate periodically, the system will quickly adapt, and as it does so the reliance upon secret knowledge techniques is replaced by a reliance upon biometrics – where the user will be continuously and non-intrusively authenticated. The result is a highly modular framework that can utilise a wide-range of standardised biometrics, and which is able to take advantage of the different hardware configurations of mobile devices – where a combination of cameras, microphones, keypads etc can be found.

When considering the hardware and form factor of a mobile device, a number of biometric techniques are found to be more applicable for deployment than others. For instance, in its present form it would not be possible to deploy a hand geometry technique as the equipment used to create the image is bulky, expensive and requires the hand to be spread flat on a surface rather than simply to be holding a device. However, various other options could be viable, and Figure 1 illustrates a number of biometrics that would (in principle) be applicable to a mobile handset.

[1] The NICA framework is based upon prior work undertaken by the authors and builds upon research originally published as the IAMS architecture (Clarke & Furnell, 2007).

Figure 1: Biometrics applicable for mobile devices

The inclusion of a camera for video calling – a standard service for third generation networks – would permit the use of facial recognition. Given sufficient picture clarity, iris scanning could also be utilised. The microphone, present for telephony services, would open the potential for voice verification, and the keypad would allow a keystroke analysis technique to be applied. For handsets or PDAs without a keypad, a touch sensitive screen is usually provided as the human-computer interface, where signature recognition could subsequently be utilised.

In practice however, all of these techniques do not currently have the functionality to be deployed in this manner, each requiring varying degrees of modification or development. Keystroke analysis, although commercially available for static-based authentication on PC keyboards, currently has no dynamic-based approach – although this technique has been thoroughly researched (Leggett et al., 1991; Napier et al., 1995). Of more concern is the applicability of keystroke analysis on a mobile handset or PDA, where the keypad or thumb sized keyboard represents a different tactile environment with which the user must interact. Preliminary studies by the author have supported this (Clarke et al., 2003; Karatzouni et al., 2007). Signature recognition has been developed commercially to provide intrusive authentication of the user based upon a signature, but not on general words signed through transcriber; although a prior study by the authors has also evaluated this (Clarke and Mekala, 2007). Speaker verification has also been developed for static (and pseudo-dynamic) authentication, but does not currently perform dynamic authentication of the user. It is clear therefore, that the majority of techniques require at least adaptation, if not a complete feasibility study before practical implementation of the technique can occur.

Research by the authors is currently underway looking at the application issues of many of these biometric techniques. It is the focus of this paper to address the applicability of facial recognition to a mobile device.

3. Facial Recognition

The use of facial recognition to date has typically focussed upon very well defined environments, with controls or restrictions placed upon the illumination, facial orientation and distance from the capture device. In a mobile device these conditions are far more variable, with authentication needing to take place under a wide-variety of different environmental conditions. The implementation of the technique in a transparent fashion will only serve to complicate these requirements further. The user will not be explicitly asked to pose as the sample is captured and could therefore suffer from a number of bad variables, such as poor lighting due to time of day or location, or having a significant difference in facial orientation as the user is looking away from the mobile device.

In order to address the issue of transparency, and thus improve the tolerance of the technique to variations, two options are available: firstly, to undertake research looking into improving the classification algorithms with a view to removing the dependence upon these factors; secondly, look to adapt current classification algorithms in a fashion that achieves transparency. This research adopts the latter choice, as research into improving classification algorithms has and will continue to take place, and designing a process that adapts existing approaches (rather than designing a single mechanism) provides more flexibility. Unfortunately, when looking to adapt current algorithms, the process is essentially balancing the FAR and FRR of the system: typically trading less security (higher FAR) in favour of a higher level of robustness and user acceptance (lower FRR).

The proposed method of adapting existing algorithms is to move away from a one-to-one comparison of an image with a template (as depicted in Figure 2(a)), and replace the template with a series of images that represent various facial orientations of the authorised user (as illustrated in Figure 2(b)). In this way, existing pattern classification algorithms can still be applied, but the overall approach should be more resilient to changes in facial orientation. Under this proposed mechanism, each sample will effectively be compared to a series of images stored within the composite template and the number of verifications will subsequently increase. This will therefore introduce an increased likelihood that an impostor is accepted by an appropriate similarity with at least one of the series of images. Under this proposed system, the FAR will only ever be as good as the original FAR of the algorithm being used, with more realistically an increase in the FAR being experienced. Conversely however, under this proposed system the FRR will at worst equal that of the previous FRR, but more realistically will be lower.

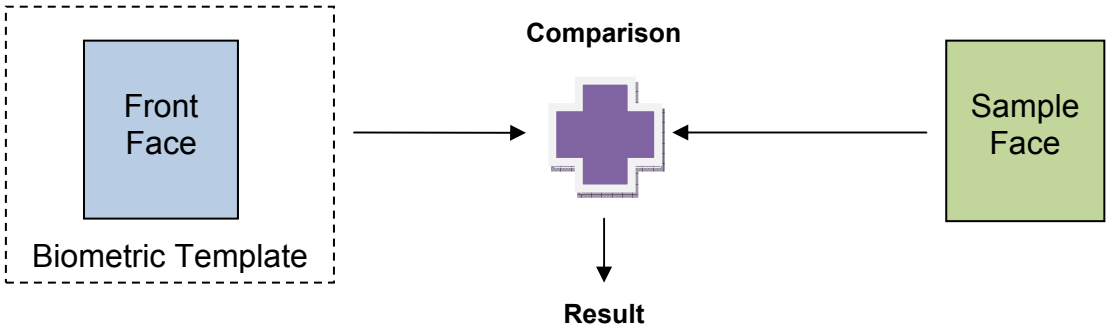

Figure 2(a): Normal facial recognition process

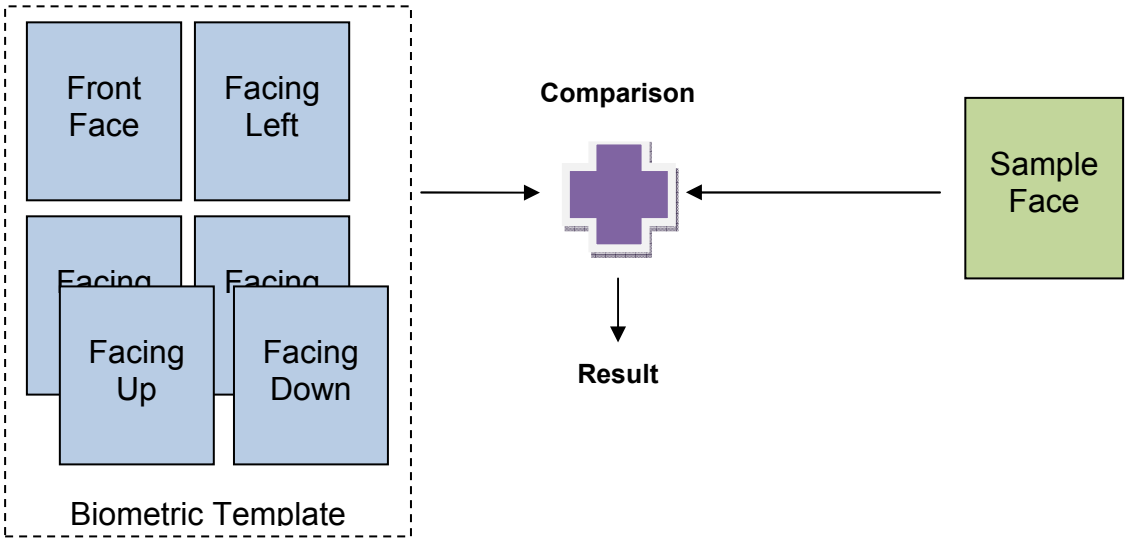

Figure 2(b): Proposed facial recognition process

The advantage of trading of the FAR and FRR in facial recognition is two-fold:

1. Facial recognition approaches have quite distinctive characteristics and experience good levels of performance in terms of FAR and FRR. Indeed, facial recognition systems are often used in identification systems as well as verification systems. The use of them for verification does not require such a high degree of distinctiveness and can therefore be traded-off with usability.
2. The relationship between the FAR and FRR is non-linear, with small changes in the FAR possibly resulting in larger changes in the FRR.

It is therefore hypothesised that it is possible to take advantage of these properties to provide a little less security for a larger improvement in the robustness and usability of the approach. The experimental study was developed in order to assess the trade-off between the error rates.

4. Experimental Methodology

The purpose of this experiment is to evaluate the change in performance rates experienced when using a composite facial template versus the traditional single front facing image template. In order to facilitate this study a series of experiments were devised to test various aspects of the hypothesis:

1. A control experiment where the facial recognition system would be tested under normal conditions.
2. An experiment to evaluate the effect upon the performance rates when using images of varying facial orientation against a normal template.
3. An experiment to evaluate the effect upon the performance rates when using images of varying facial orientation against the proposed composite template.
4. A repeat of experiment 3, but removing users that appear in the enrolment phase from the test dataset, thereby mitigating against any possible skew in the results.

In addition, it was also decided to evaluate a series of facial recognition algorithms with a view of identifying the most effective algorithm given the experimental criteria.

It was anticipated that some algorithms would be more tolerate of changes in facial orientations than others. The algorithms selected represent a number of established facial recognition algorithms. The algorithms themselves were obtained from Advanced Source Code (Rosa, 2008):

1. Fourier-Bessel Transform
2. Eigenfaces
3. Fisherfaces
4. Fourier Spectra for Faces
5. Gabor Filters

Each of the algorithms, although developed by Rosa, have been designed based upon prior published papers. For more detailed information on the algorithms themselves, please refer to Rosa (2008).

A wide variety of facial datasets exist including, the FERET colour dataset, the YALE dataset, PIE dataset, AT&T dataset, MIT dataset and NIST Mugshot Identification dataset to name but a few (Gross, 2005). All of the datasets vary in terms of the number of participants and conditions under which the images were taken – for instance, differing orientations, distance from camera, camera resolution and illumination. It was decided to utilise the FERET colour dataset as this has traditionally been a popular dataset for use within experimental studies and importantly contains the varying facial orientations required for this set of experiments.

The FERET dataset is one of the largest facial datasets with over 14,000 images. For these experiments a sub-set of the dataset was utilised and is illustrated in Table 1. Each set contains 200 images from 200 participants, with each set containing the same participants.

Dataset Ref	Description	Angle	FERET Ref
1	Front Face	0	ba
2	Alternative Front Face	0	bj
3	Left Image	+60	bb
4	Left Image	+40	bc
5	Left Image	+25	bd
6	Left Image	+15	be
7	Right Image	-15	bf
8	Right Image	-25	bg
9	Right Image	-40	bh
10	Right Image	-60	bi

Table 1: Subset of the FERET dataset utilised in the experiments

In all bar the fourth experiment, each of the participants took a turn playing the role of the authorised user, with all remaining users acting as impostors. The images used for creating the biometric template, whether it is the normal or proposed composite template, are not used in the evaluation phase – this includes experiment 4. Although this approach to biometric evaluation is standard practice, it does raise the possibility that the results produced are mildly skewed towards a better performance than can actually be produced. This is because although no image is

used in both the enrolment and verification, the algorithms have been specifically taught to reject the impostor; so should therefore be better equipped to do so over a completely new user whose sample(s) do not appear in the enrolment phase at all. The final experiment has therefore been designed to evaluate the performance of the approach with a different group of users participating in the enrolment and verification stages. The specific details of which FERET datasets were utilised in each of the experiments is illustrated in Table 2.

Exp	Enrolment		Verification	
	Dataset Ref	# of Participants	Dataset Ref	# of Participants
1	1	200	2	200
2	1	200	2, 3, 4, 5, 6, 7, 8, 9, 10	200
3	1, 3, 5, 8, 10	200	2, 4, 6, 7, 9	200
4	1, 3, 5, 8, 10	50	2, 4, 6, 7, 9	150

Table 2: Breakdown of datasets utilised in each experiment

The experiments were carried out using MathWorks MatLab to process the images, perform enrolment and verification and finally calculate the resulting false acceptance and false rejection rates (referred to as FAR and FRR respectively in the results section).

5. Results

The underlying performance of the five algorithms, as outlined by experiment 1 and illustrated in Table 3, suggests each of the algorithms are effective in rejecting impostors with a FAR of 0.19% or below. Unfortunately, the accompanying FRR is considerably larger with error rates between 21-39% for four of the techniques. The only technique that performed well was Gabor Filters with a FRR of 4.5%. It should be noted that the larger values of FRR could be as a result of fewer actual verification samples as compared to the impostors. For instance, for each user, the FRR is based upon a single image, resulting in a FRR of either 0% or 100%; whereas the FAR is based upon 199 other images, resulting in the FAR increasing in steps of 0.5%. This is unfortunately a result of the lack of repeated images per user for each orientation in the dataset. Nevertheless, the purpose of this experiment was to largely understand and establish the level of security rather usability of the underlying classification algorithms.

Algorithm	FRR (%)	FAR (%)
Fourier-Bessel	31.5	0.16
Eigenfaces	39	0.19
Fisherfaces	21	0.11
Fourier-Spectra	24.5	0.12
Gabor Filters	4.5	0.023

Table 3: Experiment 1 results

When applying the additional facial orientations to the verification process, increasing the number of images utilised to calculate the FRR from 1 to 5 (and

thereby mitigate some of the effect that a lack of comparisons can cause) the FRR increases across all algorithms. Even the Gabor-Filters approach that achieved a 4.5% FRR in experiment 1 has now increased to 46.2%. The FAR has marginally increased across the five algorithms, but is still able to provide a good level of security against impostors.

Algorithm	FRR (%)	FAR (%)
Fourier-Bessel	50.8	0.25
Eigenfaces	48.2	0.24
Fisherfaces	31.8	0.16
Fourier-Spectra	38.3	0.19
Gabor Filters	46.2	0.23

Table 4: Experiment 2 overall results

The results from this experiment demonstrate the inability of current facial recognition algorithms to cope with input samples that have a high degree of variability in facial orientation. Indeed, with current levels in the FRR, none of the algorithms evaluated in this experiment would be of any practical relevance.

Analysing the results from experiment 2 in more depth, it becomes apparent where the majority of errors reside. Figure 3 and Figure 4 present the FRR and FAR respectively as the angle of the facial orientation varies. It is clear from Figure 3 that as the angle of orientation increases an increase in the FRR is also experienced. This relationship is expected as the template is generated upon a front facing (zero degree) image which will differ to a larger degree the more obtuse the angle. It is interesting to note that no such relationship exists in the FAR, with the performance broadly flat across each facial orientation.

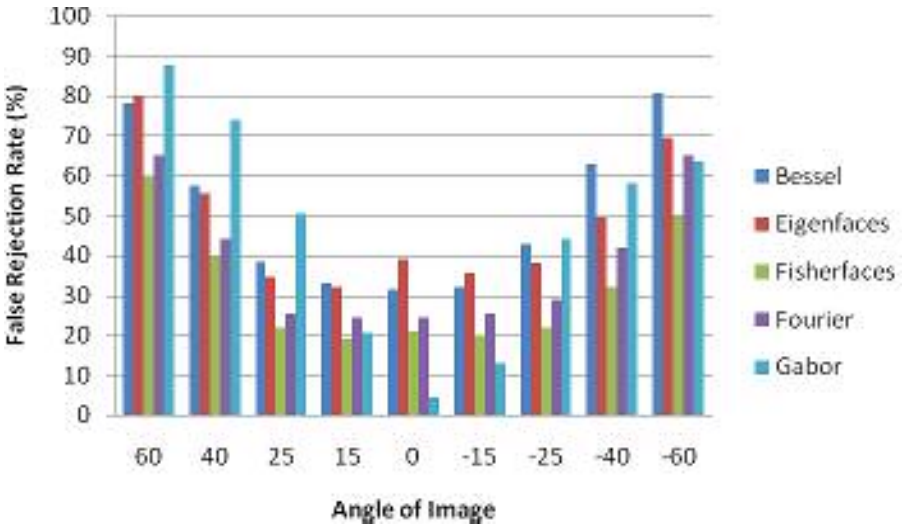

Figure 3: The affect upon the FRR with varying facial orientations

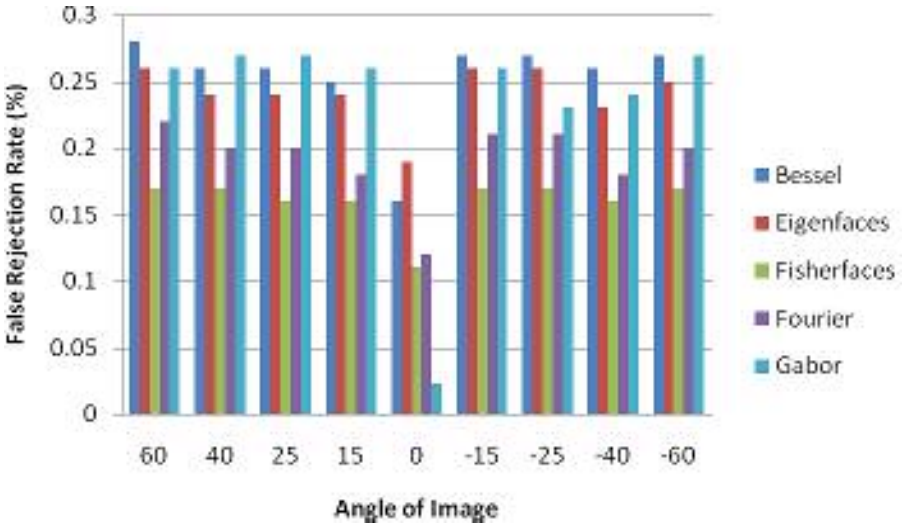

Figure 4: The affect upon the FAR with varying facial orientations

The overall results for experiment 3 are very encouraging, with the FRR having reduced considerable from the previous experiment. Indeed, the FRR are all lower than in experiment 1, demonstrating that the usability of the underlying algorithm can be improved through the use of a composite template. The important consideration is what effect this improvement has upon the level of security being provided. As illustrated in Table 5, the FARs have also improved when compared to experiment 1 (Table 3). That said, care should be taken when interpreting these results. As indicated in the methodology, theory shows that the FAR in this experiment should be equal to or larger than the standard FAR in experiment 1. In this particular experiment, this is not the case as the type and number of verifications performed differs. Nevertheless, an important observation from this data is that the FAR has only marginally changed with an accompanying large reduction in the FRR. Unfortunately it was not possible to calculate the results for the Eigenface algorithm as it proved too computationally intensive for the evaluation machine.

Algorithm	**FRR (%)**	**FAR (%)**
Fourier-Bessel	7.8	0.04
Fisherfaces	1.1	0.006
Fourier-Spectra	3.8	0.02
Gabor Filters	0.6	0.003

Table 5: Experiment 3 overall results

An analysis of the performance against the angle of facial orientation, as illustrated in Figure 5 and Figure 6, shows that the composite template is now far better placed to successfully verify users with varying degrees of facial orientations. Indeed, the worst performing orientation from the results is the traditional front-face (zero degrees) image.

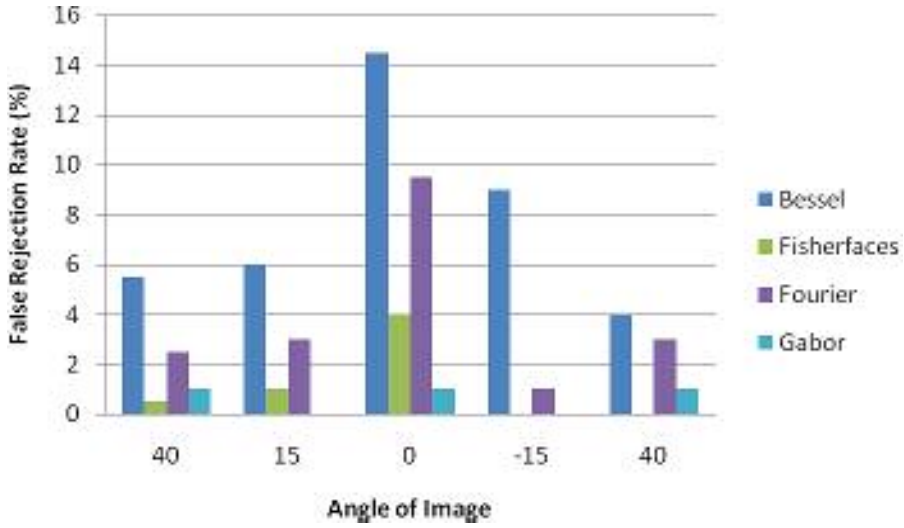

Figure 5: The affect upon the FRR using a composite facial template

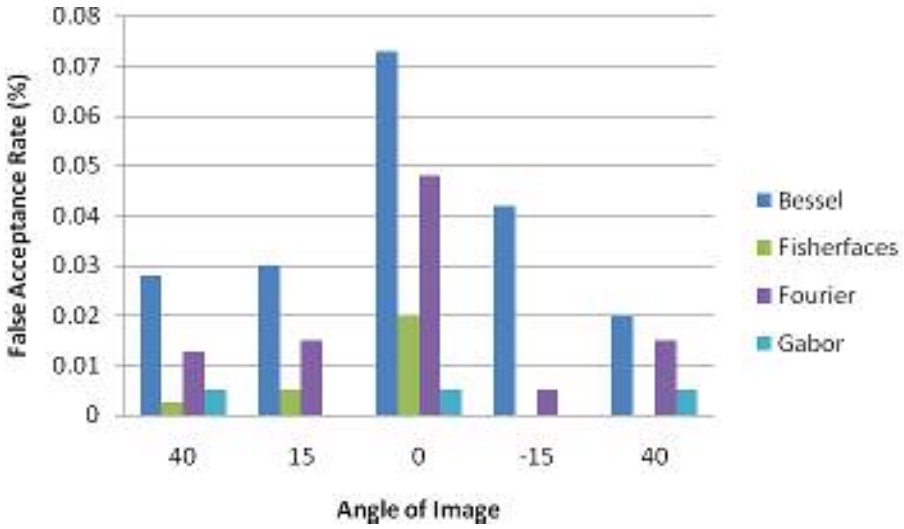

Figure 6: The affect upon the FAR using a composite facial template

Interestingly, it can be observed that a close relationship exists between the algorithms across all facial orientations; an algorithm performing well in one orientation appears to be successful across all in terms of both the FRR and FAR. This suggests that none of the algorithms are useful in classifying a particular facial orientation successfully. However, it is clear that overall the Gabor Filters and Fisherface algorithms are the best performing.

The final experiment sought to mitigate against any possible distortion of the results by using the same group of users acting as impostors in enrolment and verification. As illustrated in Table 6, any skewing that has occurred in the previous experiments has had minimal effect, with the FRR and FAR still performing well under these conditions.

Algorithm	FRR (%)	FAR (%)
Fourier-Bessel	2.8	1.87
Eigenfaces	0.8	1.87
Fisherfaces	0	1.87
Fourier-Spectra	0.4	1.87
Gabor Filters	0	1.87

Table 6: Experiment 4 overall results

It is worth highlighting, the identical FARs are a result of the averaging process that occurs against the 150 impostors. Examination of the individual FARs shows differing levels of FAR for different users in different algorithms.

6. Discussion & Conclusions

The experimental results have illustrated that utilising a composite template consisting of numerous facial orientations does improve the overall usability of the system without worsening the level of security being provided. Although the results have shown even an improvement in the FAR experienced, this is unlikely in practice as the nature of the composite template gives rise to a larger number of verification attempts and therefore an increased probability that an impostor will get access. Nevertheless, the fact the FAR was not inadvertently effected with this dataset does demonstrate that a significant improvement in the FRR can be experienced with only a marginal change in the FAR.

The actual performance of the five algorithms varied considerably, with Gabor Filters performing the best overall and Fisherfaces the second. All five algorithms were selected due to their relatively fast computation versus neural network based approaches that required intensive training periods. As such, either of the algorithms would be suitable for deployment within NICA, albeit not necessarily on the device, as this would depend upon the processing capabilities of individual devices.

The experiments presented in this study have focussed upon improving the usability of facial recognition algorithms when faced with varying facial orientations; a serious issue when looking to deploy this technique transparently. Unfortunately, this study, due to limitations in the datasets available, has not been able to test other factors that are considered essential when looking to deploy this technique to mobile devices. Factors such as camera resolution, distance between the camera and face and illumination are all key factors that vary the performance of current algorithms and future research needs to focus upon the impact these factors will have upon the composite template and the resulting performance rates.

From a wider perspective, research is still continuing into the applicability of other biometric techniques, such as signature recognition, voice verification, keystroke analysis and service utilisation. The authors are also currently developing the over-arching framework that will support the use of these techniques in a flexible and transparent manner.

Acknowledgement

This research was support by a two year grant from the Eduserv Foundation entitled "Flexible and non-intrusive user authentication for mobile devices". For more information on this research and other research outputs, please refer to www.cisnr.org/nica.

References

Ashbourn, J. (2000). Biometrics: Advanced Identity Verification. Springer

Best, J. (2006): "3G reaches 50 million users worldwide", http://news.cnet.co.uk/mobiles/0,39029678,49251672,00.htm

Clarke N, Furnell S, Lines B, Reynolds P. (2003). "Using Keystroke Analysis as a mechanism for Subscriber Authentication on Mobile Handsets". Proceedings of the IFIP SEC 2003 Conference, Athens, Greece, 26-28 May, pp. 97-108.

Clarke, N.L., Furnell, S.M. (2005): "Authentication of Users on Mobile Telephone - A Survey of Attitudes and Practices", Computers & Security, vol. 24, no.7, pp.519 - 527.

Clarke, N.L., Mekala, A.R. (2007). "The Application of Signature Recognition to Transparent Handwriting Verification for Mobile Devices". *Information Management and Computer Security*, vol. 15, no. 3, pp. 214-225.

Karatzouni, S., Clarke, N.L., Furnell, S.M. (2007). "Keystroke Analysis for Thumb-based Keyboards on Mobile Devices", Proceedings of the 22nd IFIP Information Security Conference (IFIP SEC 2007), South Africa, 14-16 May, pp. 253-263.

GSM World. (2008). "GSM Subscriber Statistics". GSMWorld.Com. http://www.gsmworld.com/

Gross, R., Face Databases, Handbook of Face Recognition, Stan Z. Li and Anil K. Jain, ed., Springer-Verlag, February 2005.

Leggett, J., Williams, G., Usnick, M. (1991). "Dynamic Identity Verification via Keystroke Characteristics". International Journal of Man-Machine Studies, vol. 35, issue. 6, pp. 859-870.

Napier, R., Laverty, W., Mahar, D., Henderson, R., Hiron, M., Wagner, M. (1995). "Keyboard User Verification: Toward an Accurate, Efficient and Ecological Valid Algorithm". International Journal of Human-Computer Studies, vol. 43, pp213-222.

Rosa, L. (2008). "Biometric Source Code". Advanced Source Code. http://www.advancedsourcecode.com/ [Accessed: 18th Feb 2008].

Smith, R. (2002). Authentication: From Passwords to Public Keys. Addison Wesley.