Face Verification System Architecture Using Smart Cards

Thirimachos Bourlai Kieron Messer Josef Kittler
University of Surrey, CVSSP Group, Guildford, GU2 7XH, U.K
{t.bourlai, k.messer, j.kittler}@surrey.ac.uk

Abstract

A smart card based face verification system is proposed in which the feature extraction and decision making is performed on the card. Such an architecture has many privacy and security benefits. As smart cards are limited computational platforms, the face verification algorithms have to be adapted to limit the facial image representations. This minimises the information needed to be sent to the card and lessens the computational load of the template matching. Studies performed on the BANCA and XM2VTS databases demonstrate that by limiting these representations the verification performance of the system is not degraded and that the proposed architecture is a viable one.

1. Introduction

Automatic personal identity verification systems based on facial images [2, 7, 8] have many promising applications in the field of security. They have the potential to confirm a person's identity in an effective and inherently more reliable way than standard pin codes. The higher level of security afforded by such systems is achieved by exploiting user specific biometric characteristics that cannot be easily misappropriated by impostors.

Face biometric systems comprise a camera linked to image processing and decision making subsystems (normally implemented using a general purpose computing engine). The role of the image processing stage is to detect the face in the data acquired by the camera at the time of access, as well as to perform its geometric and photometric normalisation. Features important for discrimination are then extracted from the registered and normalised image and fed into the decision making stage, the task of which is to accept or reject the claimed identity. The decision making process involves a comparison of a biometric face template, stored during the client's enrolment, with the extracted features.

The conventional architecture of a face biometric system consists of a camera located at a point of access. The requisite biometric data captured by the sensor is then transmitted to a central computer where all the image processing and decision making takes place. All biometric templates are stored centrally and the result of the verification process is then communicated back to the physical or electronic device which grants access to the successful claimant.

Although the centralised architecture is perfectly acceptable for some applications, in many scenarios it raises privacy issues which may compromise user acceptability. Also, there may be a problem of security if the system decision is transmitted over a public communication channel. For these reasons, there is a lot of interest in architectures where the biometric template database is distributed, and the processing system is located at the access point. A favoured set up is to store the biometric template on a smart card, together with a pin code defining the claimed identity. This alleviates the privacy problem. By locating the processing system close to the camera, the overall security of the architecture is improved. During access, the pin code and the biometric template are communicated to the local processor, after the smart card system has been duly authenticated.

In this paper we revisit the privacy and security concerns and propose a novel distributed architecture where some of the face image processing and decision making is carried out on the smart card itself. This avoids the problem of the biometric template and pin-code ever having to leave the card. In our system the feature extraction (which is invariably proprietary) is performed on the smart card, enhancing the overall system security even further.

By porting a part of the authentication process onto a smart card, one has to face the severe constraints and limitations that small computing platforms often impose. These may have a serious impact on the applicability and performance of the candidate face authentication algorithms that might be considered for the system

realisation. Unfortunately, the computational requirements of the available face verification algorithms[9, 5] are rarely discussed. We address this issue in relation to the algorithm adopted for smart card implementation. By optimising the parameters of the algorithm (including image spatial and grey level resolution) as well as the representation of numbers used by the smart card processor, we establish an acceptable trade-off between performance and computational complexity of the resulting system.

The rest of the paper is organised as follows. In the next section, the conventional and the proposed smart card face verification system architecture will be presented. In Section 3 the smart card used in our investigation is detailed and its limitations discussed. In Section 4 the experiments made to optimise our system are described. Finally, some conclusions are made.

2. System Architecture

In any face verification system the user must make an identity claim, usually by use of a token. In the experiments described in this paper, the token was stored on a smart card. To make a claim, the user presents himself/herself to a camera and places his/her card in the card reader. The token is read off the card and the relevant biometric template retrieved. A match between the template and the acquired image is then made.

Prior to this the user would have had to have gone through an enrolment process where their facial biometric template was created and stored in a database and/or on the smart card.

The acquired image will typically have to pass through several processing steps before the final matching takes place. These are: face detection/localisation; geometric and photometric normalisation and feature extraction.

There are two standard architectures to this system, centralised or decentralised. In a centralised architecture the image grabbed is transmitted to a central server where the biometric templates are stored in a database. All the processing and template matching is performed by the server's processors. The advantage of this architecture is that the processor capacity is high. However, the template has to be transmitted across a network, making the whole system vulnerable to attack.

In a decentralised system the biometric template is stored on the smart card. When a claim is made, the template is read off the card and all the processing and matching is done by a local processor. This has improved security advantages as the template does not have to be transmitted. However, the template must still leave the card. There is also a security problem of performing

Figure 1. The proposed smart card face verification system

the proprietary computation (feature extraction/decision making) on the local host.

A novel architecture is proposed that further enhances security and the privacy of face biometric verification systems, in which the processing is performed on the smart card storing the template. In this model the template never leaves the card and is thus much more secure. However, smart cards are limited computing devices suffering from slow channel communication, small storage and processing capability. For these reasons, it is not yet viable to do all the processing required on the card.

One could do all processing, including feature extraction, on a local host and just send the extracted features to the smart card. The card would just then be used to compare the stored template to the sent features and make the verification decision. However, there are still security issues in performing the proprietary feature extraction on a local host and the extracted features can be larger than the original image resulting in more data having to be transmitted. It also makes it impossible to update the system as any new feature extractor will render the current templates useless.

We propose to perform the feature extraction and decision making on the smart-card as the best compromise. A schematic diagram of the proposed architecture is shown in figure 1.

3. The Smart Card

The smart card used for this research was provided by Sharp Laboratories, Oxford, UK. It boasts a 13.5MHz processor, 1Mbyte of EEPROM, 8Kbytes of RAM, a cryptographic co-processor and operates in both contact and contactless modes. In order to perform the neces-

sary on-card verification experiments we used it in contact mode, which has a data transfer rate of 115.2Kbits per second. Undoubtedly, compared to modern computers smart cards today still offer a very limited computing platform. Primarily, smart cards have been designed for the sort of applications which allow storing and retrieving of personal information about an individual, i.e. name, address, bank account details. They have not been designed for doing the large amounts of processing that are required for biometric authentication. Moreover, no smart card has yet been manufactured which has a floating point co-processor. This makes complex mathematical calculations computationally expensive. Also, the transmission rate of data between the server and card is fairly slow. Therefore, the amount of data being exchanged between a smart card and a server through an interface (e.g. a biometric template or a facial image) must be kept to a minimum. Finally, the amount of RAM available on the smart card is limited, which means all the data can not be kept in memory for the calculations and the ROM must be used as a cache. Typically, reading data from the EEPROM is fairly fast but writing data is slow. Even when the technology on the smart card gets as powerful as a current personal computer, there will still be a requirement for making such algorithms as computationally cheap as possible.

4 System Optimisation

Based on the proposed system requirements and the known limitations of the smart card, experiments have been performed to help optimise our face verification system for the smart card platform. BANCA [6, 1] and XM2VTS [4] face databases and corresponding protocols were used. From the sets containing the face images, the training set is built and used to construct client models; then the evaluation set that produces client and impostor access scores (used to compute a client-specific or global threshold and determines acceptance or rejection of a person); and the test set that simulates realistic authentication tests where impostor's ID is unknown to the system. The threshold is set to satisfy certain performance levels on the evaluation set. The performance levels of the verification system are FA, FR and Half Total (HTER is equal to FA plus FR divided by two) Error Rates on the test set.

Most face verification algorithms rely on either PCA (Eigen) or LDA (Fischer) techniques. The feature extraction processor requirements with either of these techniques is too costly. For this system we used Client Specific LDA which has a computationally cheap feature extraction stage and is ideally suited to small platforms. The feature extraction is a simple vector dot

product of the template and normalised probe image. The resulting sum is then compared to a threshold which determines how close the probe of the claimed ID is to the class of impostors and thus the verification success. The effectiveness and viability of the method was tested in [3] where normalised correlation was used. However, in our system we are testing the impostor rejection, a formulation that obtains even better results.

In the absence of on-card floating point co-processor, all the floating point arithmetic is emulated by software on the card. The first experiment conducted was to check whether the use of fixed-point arithmetic for the feature extraction and matching step on the card could affect the verification performance in terms of speed and performance. By performing experiments using different data types on the card, it was found that the computational speed of the feature extraction and template match on the card can be increased by a factor of six when using fixed-point arithmetic over floating point numbers. In addition, as long as the position of the fixed point was chosen correctly no degradation in verification performance was observed at all.

In the proposed system the normalised images are sent to the card for verification. A limitation of the architecture is communication speed. By reducing the spatial and grey-level resolution of the normalised images, the amount of data sent to the card can be significantly reduced. In the next two experiments we investigated the effect of reducing resolution on the verification performance.

First, the grey-scale pixel resolution of the normalised probe images was altered. Since an 8 bit camera was used, the initial performance of our system was measured for the full 8 bits per pixel grey-level resolution (the spatial resolution was kept at 55x51). Then, the grey-scale resolution was reduced by 1 bpp and the performance was measured again. This was repeated down to 1 bpp. Figure 2 shows an example of an initial BANCA image and the resulted normalised bit precision images. Moreover, in figure 3 we can see typical HTER results obtained for a specific BANCA protocol. The HTER for 4 bpp was slightly less than when using 8 bpp. The amount of data sent to card could be halved without any drop in verification performance.

Next the spatial resolution of the images was varied from 110x102 down to 8x7. The grey-scale resolution was kept at 8 bpp. Table 1 shows a summary of some of the results. For the XM2VTS database the image size can be reduced (from 110x102 to 30x28) and over 13 times less data need to be sent to the smart card. Therefore, the computation for the template matching on the smart card is significantly reduced while the performance is slightly increased. In the case of the BANCA

Figure 2. Initial BANCA (first to the left) and normalised bit precision images, 1-bit to 8-bit.

Figure 3. Effect of bit precision on protocol G of the BANCA database.

database the largest size produces the best results. This is because the BANCA data has a more challenging operational scenario. However, using a size of 61x57 results in over three times less data being transmitted to the smart card with only a small decrease in performance.

Image Resolution	BANCA <i>HTER</i>	XM2VTS <i>HTER</i>
8x7	0.149	0.0865
30x28	0.107	0.03475
61x57	0.0753	0.0369
110x102	0.0664	0.03864

Table 1. Effect of image size on HTER on BANCA and XM2VTS databases.

5. Discussion and Conclusions

A smart card based face verification system has been proposed in which the feature extraction and template matching is performed on the card. This architecture offers increased security and privacy in comparison with conventional architectures. The solution involves a small footprint computing platform with severe resource limitations in terms of memory, computing engine and communication channel capacity. Consequently, the porting of algorithms on such an architec-

ture requires a careful optimisation of the system design parameters, including fixed point data types, grey-level and spatial image resolution.

Experiments showed that in order to optimise the smart card design, fixed-point arithmetic can be used which speeds up the template matching on the card. Using less than 8 bpp grey-scale image resolution for the normalised face images does not necessarily result in a degradation of the system performance. This allows for fewer bytes of data to be sent to the smart card. There is a trade-off between performance and image resolution that can be exploited to minimise data transfer and processor load. However, different results were achieved on the BANCA database in contrast to the XM2VTS one suggesting that different operational scenarios may call for different optimum operational point settings.

Acknowledgements

The authors would like to acknowledge the support received from OmniPerception Ltd, Sharp Laboratories, U.K and EPSRC under grant GR/S46543/01(P).

References

- [1] E. Bailly-Baillire, S. Bengio, F. Bimbot, M. Hamouz, J. Kittler, J. Mariethoz, J. Matas, K. Messer, V. Popovici, F. Poree, B. Ruiz, and J.-P. Thiran. The banca database and evaluation protocol. *AVBRA*, 2003.
- [2] J. L. Dugelay, J. C. Junqua, C. Kotropoulos, R. Kuhn, F. Perronnin, and I. Pitas. Recent advances in biometric person authentication. *ICASSP (special session on biometrics)*, Orlando, Florida, May 2002.
- [3] K. Messer, J. Kittler, M. Sedaghi, S. Marcel, C. Marcel, S. Bengio, F. Cardinaux, C. Sanderson, J. Cysz, L. Vandendorpe, S. Srisuk, M. Petrou, W. Kurutach, A. Kadyrov, R. Paredes, B. Kepenekci, F. B. Tek, G. B. Akar, and F. and N. Mavity. Face verification competition on the xm2vts database. *AVBRA*, pages 964–974, 2003.
- [4] K. Messer, J. Matas, J. Kittler, J. Luettin, and G. Maitre. Xm2vtsdb: The extended m2vts database. *AVBRA*, pages 72–77, March 1999.
- [5] P. J. Phillips, M. McCabe, and R. Chellappa. Biometric image processing and recognition. *EUSIPCO*, 1998.
- [6] M. Sadeghi, J. Kittler, A. Kostin, and K. Messer. A comparative study of automatic face verification algorithms on the banca database. *AVBRA*, pages 35–43, 2003.
- [7] R. Sanchez-Reillo. Including biometric authentication in a smart card operating system. *AVBPA*, pages 342–347, June 2001.
- [8] R. Sanchez-Reillo, L. Mengibar-Pozo, and C. Sanchez-Avila. Microprocessor smart cards with fingerprint used authentication. *IEEE AESM*, pages 22–24, March 2003.
- [9] W. Zhao, R. Chellappa, A. Rosenfeld, and P. Phillips. Face recognition: A literature survey. *UMD CfAR Technical Report CAR-TR-948*, 2000.