

# Sample images can be independently restored from face recognition templates

Andy Adler

School of Information Technology and Engineering  
University of Ottawa, Ontario, Canada

**Running Head:** Images can be restored from face recognition templates

**Corresponding Author:** Andy Adler  
School of Information Technology and Engineering (*SITE*),  
University of Ottawa  
800 King Edward ave.  
Ottawa, Ontario, Canada, K1N 6N5  
Tel: (613) 562-5800 ext. 6218, Fax: (613) 562-5664  
Email: aadler@uottawa.ca

## **Abstract**

Biometric systems record a sample image, and calculate a template – a compact digital representation of the essential features of the image. It has traditionally been considered infeasible to regenerate the source image from the template, and, therefore, templates are currently treated as non-identifiable data. This paper describes simple algorithm which allows regeneration of a sample image from a face recognition template using only match score values. An initial estimate of the target face is (possibly arbitrarily) chosen. The algorithm then iteratively improves this estimate to better match the target. At each iteration, a small constant times an eigenface image is added to the estimate, and modifications which improve the match score are kept. The images calculated are of sufficient quality to: 1) masquerade to the algorithm as the target person, and 2) give a good visual impression of the person's characteristics. This algorithm is immune to template encryption; any system which allows access to match scores effectively allows sample images to be regenerated in this way. This work implies that biometric templates and biometric match scores be considered identifiable data – they should not be made available to untrusted parties.

**Keywords:** Face recognition, Image restoration, Security

# 1. Introduction

Biometric systems allow automatic identification, or identity verification, of individuals using behavioural or physiological characteristics [32]. There is increasing interest in biometric authentication, especially in the context of heightened interest in national security since the attacks of September 11, 2001. Technologies such as automatic face recognition, fingerprint and iris identification, are being piloted or implemented at airports, for government identification systems such as passports and drivers licenses, and in surveillance applications [23]. Increased use of biometrics has encouraged increasing concern about the privacy and security implications of these technologies [3]. In this paper, the identifiability of stored biometric information, and its implications for biometric privacy and security is considered.

Biometric authentication is typically performed by a sophisticated software application, which manages the user interface and database, and interacts with a vendor specific, proprietary biometric algorithm. Algorithms undertake the following processing steps: 1) acquisition of a biometric sample image, 2) conversion of the sample image to a biometric template, 3) comparison of the new (or "live") template to previously stored templates, to calculate a match (or similarity) score [33]. High match scores indicate a likelihood that the corresponding images are from the same individual. The template is a (typically vendor specific) compact digital representation of the essential features of the sample image. Biometric algorithm vendors have uniformly claimed that it is impossible or infeasible to recreate the image from the templates [4,5,11,18]. These claims are supported by: 1) the template records features (such as fingerprint minutiae) and not image primitives, 2) templates are typically calculated using only a small portion of the image, 3) templates are small – a few hundred bytes – much smaller than the

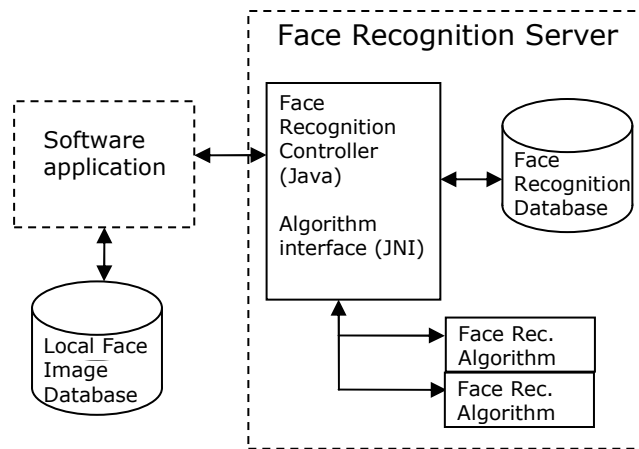
sample image, and 4) the proprietary nature of the storage format makes templates infeasible to "hack". For these reasons, biometric templates are considered to be effectively non-identifiable data, much like a password hash [18].

Because biometric data is considered to be non-identifiable, it can be managed in ways that the source images cannot. For example, templates stored on government identification documents are exchanged between governments. This type of exchange typically requires approval by a national privacy commissioner, and the primary justification is typically the assumed non-identifiable nature of biometric templates [27]. This assumed non-identifiability is also used to allay concerns expressed by citizens and employees that their fingerprint, face, and iris images may be accessed from their storage on identification cards.

This paper demonstrates that important information about biometric source images can be retrieved from templates. An algorithm is developed which can automatically regenerate a good likeness of an original image from a face recognition template using only match score data. The possibility of regenerating identifiable images in this way has important privacy and security implications for the use and storage of biometric data. This work extends a previous conference publication on the same subject [2].

## 2. Methods

A software application was implemented with the goal of recreating a face image of a target person in a face recognition database. The application has access to a local database of face images, and has network access to a face recognition server (FRS). The software begins with only the ID of the target in the database, and is able to obtain match scores of the target compared to chosen images. Beginning with an initial face image estimate, the software gradually improves this image until a fairly high quality likeness of the target is obtained. Figure 1 shows a block diagram of the software application and the software components (FRS, local database) with which it interacts.



**Figure 1:** Software architecture. The software application accesses a local database of images, and has network access to a face recognition server (FRS). The FRS allows a uniform API to access different biometric algorithms. The FRS is a SOAP application server implemented in JAVA, where the interface to each biometric algorithm uses the Java Native Interface.

## 2.1 Face Recognition Server

The FRS is a software application framework designed to permit network access to different biometric algorithms using a uniform API; it is described in detail in [1]. Briefly, the FRS is a SOAP [29] application server implemented in JAVA, and runs in the context of the Apache Tomcat JSP environment on Windows 2000. API requests (such as upload image, create template, compare templates) are made from the application to the FRS, which are translated into interactions with the FRS database and the individual face recognition algorithms. Since each face recognition algorithm was implemented in C/C++ and provided as a Win32 static or dynamic library, it was necessary to implement the interface from the FRS to the algorithm using the Java Native Interface [19]. From the application's point of view, the FRS allows calculation of a match score between an arbitrary image and a target image in the FRS database. This API is represented in subsequent pseudocode by: `match_score= req_match_score(Image, id)` A block diagram of the FRS, including the SOAP application server and interface to face recognition algorithms is shown in figure 1.

## 2.2 Image Regeneration Algorithm

Image regeneration involves preprocessing, initial image selection, and image estimate improvement. The pseudocode of this algorithm is shown in figure 2. *Preprocessing*: A local database of frontal pose face images is obtained; it does not need to resemble the target image, and may be one of the many freely available face image databases [12,20,21,25,28]. Results in this paper were calculated using the University of Aberdeen face recognition database [9,10]. Images in the local database are rotated, scaled cropped, and histogram equalized such that all images have the same size (150×200 pixels), eye locations (horizontal, vertical pixel coordinates

of the left and right eyes of 50,90 and 100,90), and pixel intensity distribution. Then the first 400 eigenimages (also called principle components or eigenfaces [30]) were calculated, using the method of [15]. *Initial image selection:* The software determines the match score for a selection of images in the local database against the target template. The initial image estimate is selected to be the one with the highest match score. In this paper, initial image estimates were chosen to be visually dissimilar from the target, in order illustrate the effect of optimization. *Image estimate improvement:* After the initial estimate is selected, it is modified by the algorithm to better match the target. In each iteration, an eigenimage is selected, and a series of images produced equal to the current image estimate plus a (heuristically determined) small constant ( $c$ ) times the eigenimage. The match scores between these images and the target are calculated, and the image with the best score is selected for the subsequent iteration. Eigenimages are selected in order of decreasing eigenvalue, repeating with the first eigenimage after the 400th. However, other selection criterion (including random) produce similar results. Iterations are repeated until there is no significant improvement in match score. It was heuristically determined that six different adjustment levels for each eigenimage gave the fastest convergence. Typically, the algorithm reached a maximum match score after about 4000 iterations (or  $6 \times 4000 = 24\ 000$  match score calculations).

```

• Given
  Target ID in FR database: TID

• Preprocessing
  Normalize, center, equalize local image database:  $Img[i]$ 
  Calculate eigenface representation:  $EF[k]$ ,  $1 \leq k \leq 400$ 

• Determine starting image,  $Img[0]$ 
  for  $i = 1$  to number_images
    match_score[ $i$ ] = req_match_score(  $Img[i]$ , TID)
   $Img[0] = Img[ match\_score == \min(match\_score) ]$ ;

• Optimize image estimate,  $Img[j]$ 
  for  $j = 0$  to optimization_tries
     $EF_j = EF[ j \% \text{number\_eigenfaces} ]$ 
    for  $k = -3$  to 3
      score[ $k$ ] = req_match_score(  $Img[j] + c k EF_j$ , TID)
    set  $k_{max}$  to value which maximizes score[ $k$ ]
     $Img[j+1] = Img[j] + c k_{max} EF_j$ 
    crop  $Img[j+1]$  if values outside image bounds

```

**Figure 2:** Pseudocode of the software application for regeneration of images from face recognition templates. The steps of preprocessing, initial image selection, and image estimate improvement are outlined.

### 2.3 Match Score Normalization

Match score values calculated by different biometric algorithms cannot be (directly) compared, because the match score corresponds to an arbitrary distance metric. In order to compare the values calculated by the different face recognition algorithms tested, a methodology to convert match scores to confidence values (which can be compared) is developed. Interestingly, even though several biometrics vendors provide conversion tables between match scores and confidence values [27] there does not seem to be any public description of the calculation methodology.

Given a biometric algorithm and a sample database representative of the target population, match scores are calculated corresponding to comparisons between each pair of images in the database.

The probability distributions of match scores are calculated for genuine and impostor transactions. The distribution for impostor transactions ( $f_i(m)$ ) is calculated from all comparisons between different individuals, while the distribution for genuine transactions ( $f_g(m)$ ) represents all comparisons of the images from the same person, but not including identical images. We wish to calculate the confidence value  $c$  corresponding to a match score  $m$ . Using conditional statistical notation,  $c$  is represented as  $p(g|m)$ ; this indicates the probability that the event was genuine ( $g$ ) given a match score  $m$  was obtained. Using Bayes' rule, we can calculate:  $p(g|m)p(m) = p(m|g)p(g)$ , where  $p(g)$  is the *a priori* probability of a genuine event; it is chosen to be 50% to simplify the subsequent analysis.  $p(m|g)$  is the probability of getting match score  $m$  given a genuine event; in terms of the probability distributions,  $p(m|g)$  is  $f_g(m)$ .  $p(m)$  is the probability of  $m$  for all events, genuine and impostor; in terms of the distributions, it is  $p(g)f_g(m) + (1-p(g))f_i(m)$ . Thus:

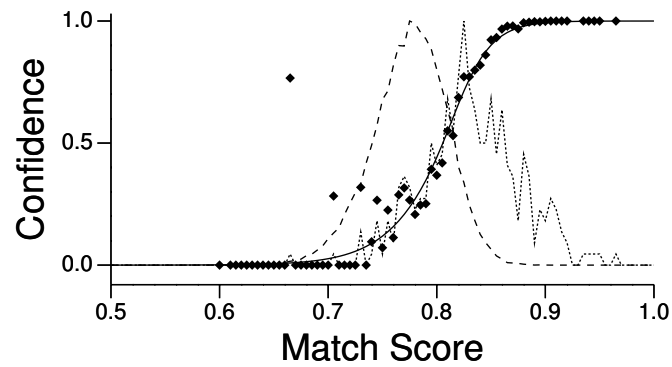
$$p(g|m) = \frac{p(g)f_g(m)}{p(g)f_g(m) - (1-p(g))f_i(m)} = \left(1 + \frac{f_i(m)}{f_g(m)}\right)^{-1} \quad (1)$$

This expression defines the conversion function between the match score and the *confidence* value. Its value depends on the sample database, and the probability of a genuine event. The general shape is sigmoid, however the curve is steeper nearer to  $c=1$  than  $c=0$ . This difference is due to outliers – genuine transactions that compare at low match score [33]. The accuracy of the calculated  $p(g|m)$  depends on the error in the distribution functions, and thus the estimated value becomes unreliable as the number of samples in  $f_i$  and  $f_g$  become small. To compensate for this estimation error, a model was fit to the data to extrapolate the overall shape of the curve. Several

variants of sigmoid functions were tested for the model; the simplest which was able to visually satisfactorily fit the data was:

$$f(m) = \frac{1}{1 + \exp\left(k_1(m - m_0) + k_2\left(\sqrt{(m - m_0)^2 + k_3^2} - k_3\right)\right)} \quad (2)$$

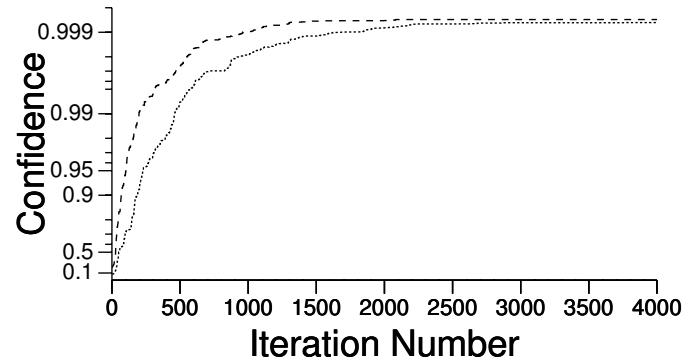
The exponential term has different asymptotic slope for values of  $m$  above and below  $m_0$ , allowing it to model the presence of outliers at low match score in the genuine distribution. Model parameter values ( $m_0, k_1, k_2, k_3$ ) were calculated from the data using a weighted fit, using the Nelder-Mead simplex function minimization algorithm. Figure 3 illustrates representative data for confidence as a function of match score. Results are shown for one face recognition algorithm using the Mugshot face database [21]. The genuine distribution appears less smooth, as it is calculated with far fewer data than the impostor distribution. Visually, the fitted curve reflects the sigmoid shape of the measured confidence values, although there is some discrepancy at low match score due to the outliers.



**Figure 3:** Confidence value as a function of match score for one face recognition algorithm. The normalized genuine (dotted line) and impostor (dashed line) distributions are shown as a function of match score, as well as the confidence value estimates ( $\diamond$ ), and the best fit model (solid line). The genuine distribution appears less smooth, as it is calculated with far fewer data than the impostor distribution. Visually, the fitted curve reflects the sigmoid shape of the measured confidence values.

### 3. Results

The image regeneration algorithm was tested using three different face recognition software packages. All are recent products by well known commercial vendors of biometric systems. Two of the vendors participated in the 2002 face recognition vendor test [13]. One target image (figure 5A and 5B lower left) was converted to (algorithm specific) templates and stored in the FRS database. Ten different initial image estimates were used. Two representative initial estimates are shown in figures 5A and 5B, upper left. As expected, initial match scores corresponded to low confidence values (on the order of 50% or less). As the algorithm calculated improved estimates, the confidence rapidly increased, reaching a maximum value after approximately 4000 iterations. In all cases, the maximum match score corresponded to a confidence above 0.999. This indicates that the face recognition algorithm would certainly accept the calculated estimate as an image of the target person. Figure 4 shows the confidence value versus iteration number for the initial estimates of figure 5A and 5B, for one face recognition algorithm. The confidence initially improves rapidly, after which improvement tends asymptotically toward its maximum.




**Figure 4:** Confidence value of regenerated image compared to the target as a function of iteration number ( $k$ ). Curves are shown for the initial estimate of figure 5A (dotted) and 5B (dashed). The vertical axis uses a logarithmic scale proportional to  $1 - \text{confidence}$ . The confidence of match between the initial image and the target (at  $k=0$ ) is low. As the algorithm iteratively improves the image estimate, confidence value increases dramatically, reaching a asymptote at above 0.999.

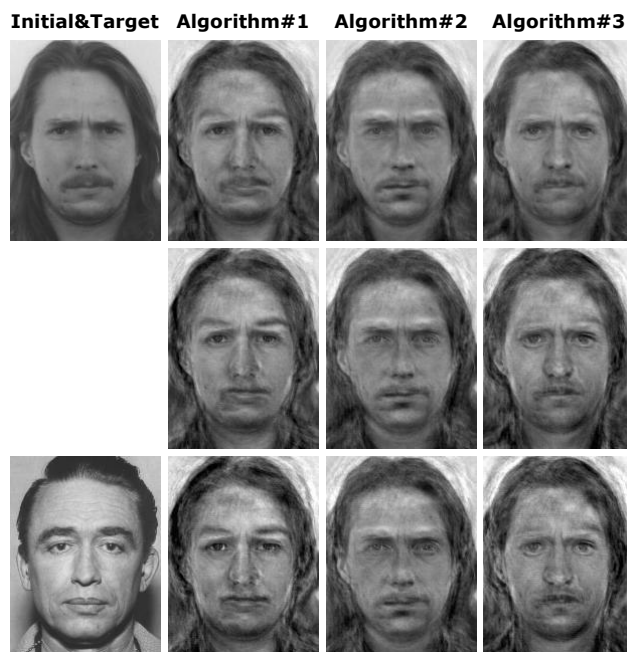
Figures 5A and 5B show representative images of the progress of image optimization. The initial image estimate and target image are shown at the top and bottom left, respectively. Each column shows a sequence of calculated image estimates for a different face recognition algorithm. From top to bottom, images are shown at iterations  $k= 200, 500,$  and  $4000$ . Corrections to the image occur primarily in the eyebrows, and shape of the eyes, nose, mouth, and upper head. The lower face shape, hair, beard/moustache region and ear shapes receive no substantial alteration, likely because this information is not encoded in the template. The shape of the mouth is modified significantly; this is somewhat surprising, as one would not expect the mouth shape to be biometrically significant. Similarly, at least for algorithm #1, the hair is quite modified – suggesting that changes in hair style may fool some face recognition algorithms. Differences in

behavior with the various algorithms are significant; for example, for algorithm #2, there is no modification in nose width or upper head shape, unlike for the other algorithms.

**Figure 5A**




**Figure 5B**



**Figure 5:** Sample regenerated images for two different initial estimates (upper left, figures 5A and 5B) for a single target (lower left, 5A and 5B). Each column shows a sequence of calculated image estimates for a different face recognition algorithm. From top to bottom, images are shown at iterations  $k=200$ , 500 and 4000. Corrections to the image occur primarily in the eyebrows, and shape of the eyes, nose, mouth, and upper head. Differences between algorithms are significant; for example, algorithm #2 does not modify the nose width.

In order to improve the visual appearance of the regenerated image, a composite image is calculated by averaging the regenerated image from ten different initial estimates. These results are shown in figure 6. The two images on the right are average estimates from face recognition algorithms one and three, while the target image is shown on the left. The software licence for algorithm #2 expired before it was possible to calculate all the images required for this figure. In all cases, the average image has a surprisingly good resemblance to the target.



**Figure 6:** Image estimates calculated by averaging the optimized results from ten different initial image estimates. The target image is shown at left, while the center and right images are average optimized images.

## 4. Discussion

Systems using biometric authentication technologies are seeing increasing use for government identification, law-enforcement, and industrial and commercial applications [13,23]. Commensurate with this usage, there is concern over the privacy and security implications, both from users whose biometric details are captured, and from the governments and industries making significant investments in biometric security systems [3]. One security and privacy factor which has previously received little attention is the identifiability of biometric templates; it has generally been assumed that it is impossible or infeasible to regenerate a biometric source image from template data. In this sense, templates are treated much like a cryptographic one way function (hash) used for password storage; it is easy to verify whether a given password corresponds to the hash, but infeasible to calculate the password from the hash [18].

In this paper, we show that, at least in the case of face recognition templates, a fairly high quality image of a person can be automatically regenerated. An algorithm is presented which can successfully regenerate images from templates from three different face recognition systems. This approach works automatically, without special user expertise; the only information required is the match score between the target template and an arbitrary image. The images calculated using the procedure are of sufficient quality to: 1) masquerade to the algorithm as an image of the target, and 2) give a good visual impression of the person's characteristics. This approach does not allow exact recreation of the target image. This is not surprising, as the template typically contains significantly less data (typically hundreds of bytes to a few kilobytes) than the original image. Conceptually, the algorithm functions by extracting information from the match score values. Since, unlike a cryptographic system, in which a passphrase is either correct or not,

a biometric algorithm provides a measure of the similarity between the sample image and the target. It is this information which allows an initial guess to be gradually refined. After comparison with many sample images, enough information is "leaked" by the match scores to permit regeneration of the target image. These results imply that both the biometric templates and match score results should be considered identifiable; these data should not be provided to untrusted parties. This work demonstrates the feasibility of this approach for face recognition systems; however, this conceptual framework should be readily extensible to other biometric modalities, as suggested by the recent work of Hill [17].

Several security concerns can be envisioned using image regeneration from biometric templates. The passage of recent U.S. legislation to require biometrics on U.S. visas and on passports from visa waiver countries [7], means that biometric data from hundreds of millions of people will be accessible in smart cards and 2D barcodes printed on these documents, from which the source images may be regenerated. Although in the case of face recognition biometrics, the privacy issue is minimal (since the holders photo is already on the document) the possibility of regeneration of fingerprint or iris images could pose significant privacy concerns. Another possible implication of this work concerns government agencies that allow collaborating agencies to perform searches against a biometric watch list. For national security reasons, the primary agency may not want to distribute watch list images, but may be effectively permitting this through image regeneration from match scores.

Previously Hill demonstrated that it was possible to reverse engineer the file format of a particular (unspecified) fingerprint algorithm [17]. Software was developed to generate an image which would compare at high match score with the original, and visually demonstrate the

primary characteristics of the original fingerprint. The implications of this work was analysed in a report by the International Biometric Group [18]. Three types of biometric image recreation were distinguished: 1) feature (an image which fools biometric algorithm, rated achievable), 2) generic image (a rough resemblance to the original, rated very likely achievable), and 3) total image (virtually identical to the original, rated very difficult, though perhaps not impossible). The biometric user was identified as being vulnerable to hostile identification and masquerade, and biometric vendors identified as potentially vulnerable to "hostile vendors". The vulnerabilities to which the institutional user of biometrics were exposed was not discussed, although this would presumably be mostly due to masquerade. Encryption of templates and trusted devices were recommended to increase protection from image recreation from templates. This paper extends Hill's results in several ways. Access to the template storage is not required. This implies that encryption of the template [18] will not protect against image regeneration if match scores data is available. Additionally, this algorithm does not require significant technical expertise to analyse and "hack" a proprietary file format [17]. As long as a uniform interface API is provided (such as by a FRS [1]), the software will execute unmodified with a different face recognition algorithm.

This algorithm is not particularly sensitive to the choice of optimization algorithm, initial image estimate, or of local face database. Tests were conducted with various initial estimates and in all cases the calculated estimate matched at high confidence to the target. Calculated estimates still retained many characteristics of the initial images, such hair and average skin tone. However, it appeared that images that were visually quite different from the target (such as different skin color) required more biometric comparisons in the optimization algorithm. The choice of optimization algorithm does not appear to be critical to this technique: although the algorithm

presented in this paper is simple, it appeared to function well. To explore the impact of the choice of optimization strategy, the algorithm was reimplemented using the Nelder-Mead simplex [16]. Results, in terms of maximum confidence values, were almost identical to the simple algorithm; however, the number of biometric comparisons required was approximately three times greater. This result suggests that there may be optimization strategies that are significantly more efficient in terms of the required number of biometric comparisons.

The choice of local face database also does not appear to be critical to the success of this method. The local face database was chosen to be the University of Aberdeen face recognition database [9,10]. This face image collection consists of largely frontal poses of primarily young caucasian university students taken with a video camera under various indoor lighting conditions; the age differences between images are small (days or weeks). In order to select target images quite dissimilar from the local face database, target images were chosen from the Mugshot database [21]. This is a collection of frontal and profile poses of primarily men taken by law enforcement officials; it is considered to be one of the more difficult for face recognition [24,31], largely because of the large age range over which different image of individuals are acquired. The most important requirement for a the algorithm is a good choice of basis functions. For face recognition, the PCA basis performs well. Several other basis functions, such as a two dimensional Fourier transform, and a pixel blocks basis were tested, with disappointing results; the algorithm produced almost no improvement in the initial estimate for these bases. For face recognition, the PCA basis is natural; many successful biometric algorithms are inspired by it [14,22,26,30]. Interestingly, none of the face recognition algorithms tested are directly based on the PCA algorithm, so the success of template regeneration does not depend on "inversion" of the underlying algorithm. These factors suggest that this method could be extended to other

biometric modalities. For this work the primary requirement would be a choice of a basis representation. For example, for fingerprint recognition, where biometric features (minutiae) are spatially localized [17], specifically optimized basis functions would be required.

Many different figures of merit have been developed to describe the performance of biometric systems. The most commonly used are parameters of false accept and false reject rates based on signal detection theory; however, many other parameters have been proposed which express biometric performance in various application scenarios [6,8,13]. None of these figures of merit were ideally suited for the analysis in this paper; in order to compare the meaning of results for biometric performance algorithms, an approach to normalize match score results was required. The approach presented allows calculation of a confidence value from match score data. This parameter may also be a useful measure on its own; in the author's experience, attempts to explain biometrics are typically confronted with the question: "But what does that match score mean?" This task is made worse by the unfortunate common usage of the term "confidence" for match score results. Some biometrics vendors publish a conversion table of match score to confidence values [27], but there has not been a description of the details of the calculation. Confidence values depend on the sample database and the probability of genuine and impostor events. This probability is application dependent; for example, in passport inspection stations it is probably quite high, while for an unsupervised building access point, it may be significantly lower. Differences in the sample database affect the calculated confidence; a "harder" database, with poorer quality images and greater time difference between captured images, will tend to reduce the match scores for genuine transactions, while having less effect on scores for impostor transactions. A "hard" sample database will thus result in larger confidence values for a given match score.

## 5. Conclusion

An algorithm has been developed by which an image of a person encoded in a face recognition biometric template may be automatically regenerated. An initial image estimate is selected and then gradually improved using the match score calculated between test image modifications and the target template. In tests with three different face recognition algorithms, regenerated images were of sufficient quality to masquerade to the algorithm as the target, and to give an overall impression of the target person's characteristics. Importantly, this approach does not require a technically sophisticated user (as did Hill [17]), and cannot be protected against by encryption of the template. The simplicity of this algorithm suggests that it be extensible to other biometric modalities. This suggests that, in terms of privacy and security of biometric systems, access to biometric templates and match score results can effectively allow access to identifiable source images. This work implies that biometric templates and match scores should be considered identifiable data – they should not be made available to untrusted parties.

## 6. References

1. A. Adler, "Automatic Face Recognition System Architecture for Collaborative Virtual Environments", *IEEE Int. Workshop Haptic Virt. Environ.*, 1:1-6, 2002.
2. A. Adler, "Sample images can be independently restored from face recognition templates", *Can. Conf. Elec. Computer Eng.*, May 2003.
3. J. Alexander, J. Smith, "Engineering Privacy in Public: Confounding Face Recognition", *Workshop on Privacy Enhancing Technologies (PET2003)*, Dresden, Germany, 2003.  
<http://petworkshop.org/preproc/07-preproc.pdf> (current June 2003)
4. H. Arendt, "Biometric Identification and National Security", *Secure*, 5:56-57, 2002.  
[http://www.silicon-trust.com/pdf/secure\\_5/56 techno\\_6.pdf](http://www.silicon-trust.com/pdf/secure_5/56 techno_6.pdf) (current June 2003)
5. Argus Solutions, *Technology Overview – Privacy*, <http://www.argus-solutions.com/Privacy.html> (current June 2003)
6. J.R. Beveridge, K. She, B. Draper, G.H. Givens, "A nonparametric statistical comparison of principal component and linear discriminant subspaces for face recognition", *Proc. IEEE Conf. Computer Vision Pattern Recog.*, pp. 535-542, 2001.
7. Congress of the U.S.A., *Enhanced Border Security and Visa Entry Reform Act*, 2002.  
[http://unitedstatesvisas.gov/pdfs/Enhanced\\_Border\\_SecurityandVisa\\_Entry.pdf](http://unitedstatesvisas.gov/pdfs/Enhanced_Border_SecurityandVisa_Entry.pdf) (current June 2003)
8. M. Bone, D. Blackburn, *Face Recognition at a Chokepoint: scenario Evaluation Results*, 2002. [http://www.dodcounterdrug.com/facialrecognition/DLs/ChokePoint\\_Results.pdf](http://www.dodcounterdrug.com/facialrecognition/DLs/ChokePoint_Results.pdf) (current June 2003)
9. I. Craw, D. Tock, A. Bennett, "Finding face features", *Proc. Euro. Conf. Computer Vision*, 588:92-96, 1995.

- <http://www.maths.abdn.ac.uk/maths/department/preprints/files/95114.ps> (current June 2003)
10. I. Craw, N.P. Costen, T. Kato, S. Akamatsu, "How should we represent faces for automatic recognition?", *IEEE Trans. Pattern Analysis and Machine Int.*, 21:725-736, 1999.
  11. Ethentica Corp., *FAQ: Biometric Template Information*,  
<http://www.ethentica.com/template.html> (current June 2003)
  12. P.J. Phillips, H. Moon, P.J. Rauss, S. Rizvi, "The FERET evaluation methodology for face recognition algorithms", *IEEE Trans. Pat. Analysis Machine Int.*, 22(10):1090-1104, 2000. <http://www.nist.gov/humanid/feret/> (current June 2003)
  13. P.J. Phillips, et al., *Facial Recognition Vendor Test 2002 Evaluation*, Feb. 2003.  
[http://www.frvt.org/DLs/FRVT\\_2002\\_Evaluation\\_Report.pdf](http://www.frvt.org/DLs/FRVT_2002_Evaluation_Report.pdf) (current June 2003)
  14. R. Gross, J. Shi, J. Cohn, "Quo vadis face recognition?", *Third Workshop Empirical Evaluation Methods Computer Vision*, Dec. 2001.
  15. P. Grother, "Software Tools for an Eigenface Implementation", *National Institute of Standards and Technology*, 2000. <http://www.nist.gov/humanid/feret/> (current June 2003)
  16. N.J. Higham, "Optimization by direct search in matrix computations", *SIAM J. Matrix Anal. Appl.*, 14(2): 317-333, 1993. <http://www.maths.man.ac.uk/~higham/mctoolbox/> (current June 2003)
  17. C.J. Hill, "Risk of Masquerade Arising from the Storage of Biometrics", *B.S. Thesis, Australian National University*, 2001. <http://chris.fornax.net/biometrics.html> (current June 2003)

18. International Biometric Group, "Generating Images from Templates", *I.B.G. White Paper*, 2002. [http://www.ibgweb.com/reports/public/reports/templates\\_images.html](http://www.ibgweb.com/reports/public/reports/templates_images.html) (current June 2003)
19. S. Liang, *Java Native Interface: Programmer's Guide and Specification*, Addison Wesley, Boston, MA, USA, 1999.
20. A.M. Martinez, R. Benavente, *The AR Face Database*, Tech. Report #24, Computer Vision Center, Campus Universitat Autònoma de Barcelona, June 1998.  
<http://rv11.ecn.purdue.edu/v1/ARdatabase/ARdatabase.html> (current June 2003)
21. NIST, *NIST Special Database 18: Mugshot Identification Database (MID)*,  
<http://www.nist.gov/srd/nistsd18.htm> (current June 2003)
22. P.J. Phillips, A. Martin, C.L. Wilson, M. Przybocki, "An Introduction to Evaluating Biometric Systems", *IEEE Computer*, 33(2):56-63, 2000.  
<http://www.frvt.org/DLs/FERET7> (current June 2003)
23. P.J. Phillips, "Human identification technical challenges", *Proc Int. Conf. Image Proc.*, 1:49-52, 2002.
24. P.J. Phillips, E.M. Newton, "Meta-analysis of face recognition algorithms", *Proc. IEEE Int Conf. Automatic Face Gesture Recog.*, 5:224-230, 2002.
25. P.J. Phillips, H. Wechsler, J. Huang, P. Rauss, "The FERET database and evaluation procedure for face-recognition algorithms", *Image and Vision Computing Journal*, 16(5):295-306, 1998.
26. A. Rukhin, P. Grother, P.J. Phillips, E. Newton, "Dependence characteristics of face recognition algorithms", *Proc. Int. Conf Pattern Recog.*, 16(2):36-39, 2002.
27. J.F. Shaw, Personal communications, Jan. 2002.

28. T. Sim, S. Baker, M. Bsat, *The CMU Pose, Illumination, and Expression (PIE) database of human faces*. Tech. report CMU-RI-TR-01-02, Robotics Institute, Carnegie Mellon University, 2001.  
[http://www.ri.cmu.edu/pub\\_files/pub2/sim\\_terence\\_2001\\_1/sim\\_terence\\_2001\\_1.pdf](http://www.ri.cmu.edu/pub_files/pub2/sim_terence_2001_1/sim_terence_2001_1.pdf)  
(current June 2003)
29. D. Box, et al., "Simple Object Access Protocol (SOAP) 1.1", *World Wide Web Consortium*, <http://www.w3.org/TR/SOAP/> (current June 2003)
30. M.A. Turk, A.P. Pentland, "Eigenfaces for recognition", *J. Cognitive Neuroscience*, 3(1):71-86, 1991.
31. F. Wallhoff, S. Muller, G. Rigoll, "Recognition of Face Profiles from the MUGSHOT Database Using a Hybrid Connectionist/HMM Approach", *IEEE Int. Conf. Acoustics Speech Signal Proc.*, Salt Lake City , Utah, July 2001
32. J.L. Wayman, *A Definition of "Biometrics"*, 2001.  
<http://www.engr.sjsu.edu/biometrics/nbtccw.pdf> (current June 2003)
33. J.L. Wayman, "Fundamentals of Biometric Authentication Technologies", *Proc. Card Tech/Secure Tech*, 1999. <http://www.engr.sjsu.edu/biometrics/nbtccw.pdf> (current June 2003)