Sample Images can be Independently Restored from Face Recognition Templates

Andy Adler
School of Information Technology and Engineering, University of Ottawa, Ontario, Canada
aadler@uottawa.ca

Abstract

Biometrics promise the ability to automatically identify individuals from reasonably easy to measure and hard to falsify characteristics. They are increasingly being investigated for use in large scale identification applications in the context of increased national security awareness. This paper addresses some of the security and privacy implications of biometric storage. Biometric systems record a sample image, and calculate a template: a compact digital representation of the essential features of the image. To compare the individuals represented by two images, the corresponding templates are compared, and a match score calculated, indicating the confidence level that the images represent the same individual. Biometrics vendors have uniformly claimed that it is impossible or infeasible to recreate an image from a template, and therefore, templates are currently treated as nonidentifiable data. We describe a simple algorithm which allows recreation of a sample image from a face recognition template using only match score values. At each iteration, a candidate image is slightly modified by an eigenface image, and modifications which improve the match score are kept. The regenerated image compares with high score to the original image, and visually shows most of the essential features. This image could thus be used to fool the algorithm as the target person, or to visually identify that individual. Importantly, 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.

1. INTRODUCTION

Biometric authentication is automatic identification, or identity verification, of individuals using behavioural and/or physiological characteristics [8]. 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. In this paper, we consider the *identifiability* of stored biometric information, and its implications for biometric privacy and security.

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 score. High match scores indicate a likelihood that the corresponding images are from the same individual. The biometric 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 template. [2, 3, 4, 7] 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 [7]. In fact, these arguments are not valid: this paper demonstrates a simple algorithm to recreate sample images from templates using only match score results.

2. METHODS

A software application (figure 1) was designed with the goal of recreating a face image of a specific person in a face recognition database. The application has local access to a database of face images, and has network access to a Face Recognition Server (FRS)

CCECE 2003 - CCGEI 2003, Montréal, May/mai 2003
0-7803-7781-8/03/$17.00 © 2003 IEEE

[1]. The FRS is a SOAP application server implemented in Java, with JNI access to individual face recognition software packages. The FRS allows a uniform network based interface to multiple face recognition algorithms. The application makes SOAP API requests to obtain the match score between an uploaded image and a specific stored record ID. This API is represented is subsequent pseudocode by:

```
match_score=
req_match_score(Image,pers_id)
```

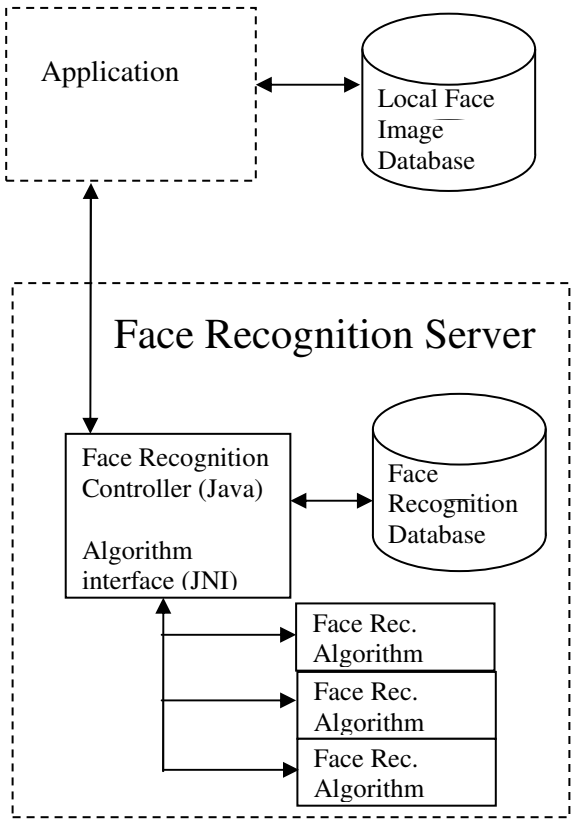

Figure 1: Software architecture of this paper. The Application accesses a local database of images, and has network access to a face recognition server.

The pseudocode of the application is shown in figure 2, involving the following steps: preprosessing, starting point determination, and image optimization. Initially, all that is available to the application is the person ID (PID) in the database of the FRS. During preprocessing, images in the local database are rotated, scaled and cropped, such that all images have the same size (150×200 pixels) and eye locations (50,90 and 100,90). Subsequently the image database was decomposed to calculate the first 300 eigenimages (also called principle components or eigenfaces) [5].

- **Given**
  Person ID in FR database: PID
- **Preprocessing**
  Normalize local image database: Img[i]
  Calculate eigenface representation: EF[k]
- **Determine starting image, Im[0]**
  for i = 0 to number_images-1
    match_score[i] = req_match_score(Img[i], PID)
    Im[0] = Img[match_score == min(match_score)];
- **Optimize image estimate, Im[k]**
  for k = 0 to optimization_tries
    curr_EF = EF[k % number_eigenfaces]
  find *c* to minimize:
      req_match_score(Img[k] + *c*×curr_EF, PID)
  Im[k+1] = Img[k] + *c*×curr_EF
  crop Im[k+1] if values outsize image bounds

Figure 2: Application Pseudocode.

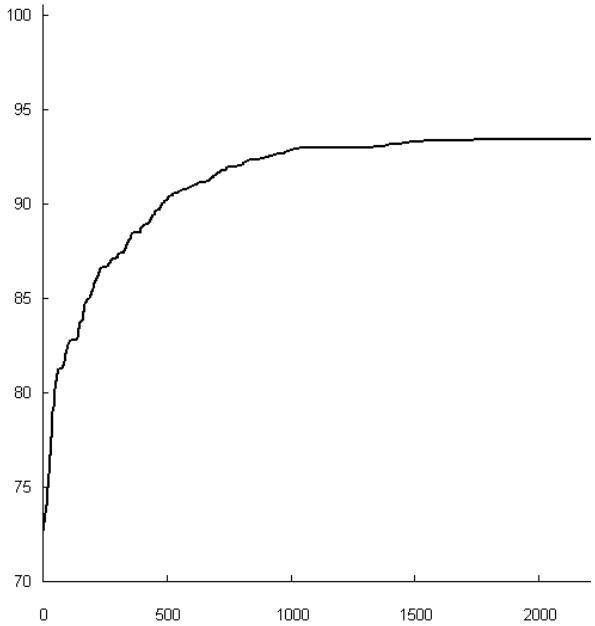

Figure 3: Match score versus number of iterations. Match score is normalized to a maximum of 100.

The application then determines the match score for all images in the local database against the PID. The initial estimate is selected to be the image with the highest match score. Subsequently, the initial estimate is improved as follows: for each iteration, an eigenimage is selected, and a series of image produced equal to the current image estimate plus a small constant times the eigenimage. The corresponding match scores are calculated, and the image with the best score is selected for the subsequent iteration. This process is repeated until there is no significant

improvement in match score. Using 6 levels at each iteration, the algorithm stabilized after about 2500 iterations.

3. RESULTS

Figure 3 shows a representative graph of the normalized match score as a function of iteration number. Most improvement occurs near the beginning, with subsequent iterations contributing little to the match score. Note that for this algorithm, the starting match score (0.72) indicates very dissimilar persons, while the maximum achieved match score (0.94) indicates a very high probability of match.

Representative images are shown in figure 4. In order to illustrate the behaviour of the algorithm, a starting image estimate (left, top) visually quite dissimilar to the target (right, bottom) was chosen. The first four images from top left are Im[k] for k= 0, 200, 500 and 3200. Corrections to the image occur primarily in the eyebrows, and shape of the nose and mouth. The face width, hair and ear shapes receive no substantial alteration, likely because this information is not encoded in the template. The fourth image from the left is an average image from four different starting images; while it does contain significant artefacts it shows the key visual characteristics of the target image.

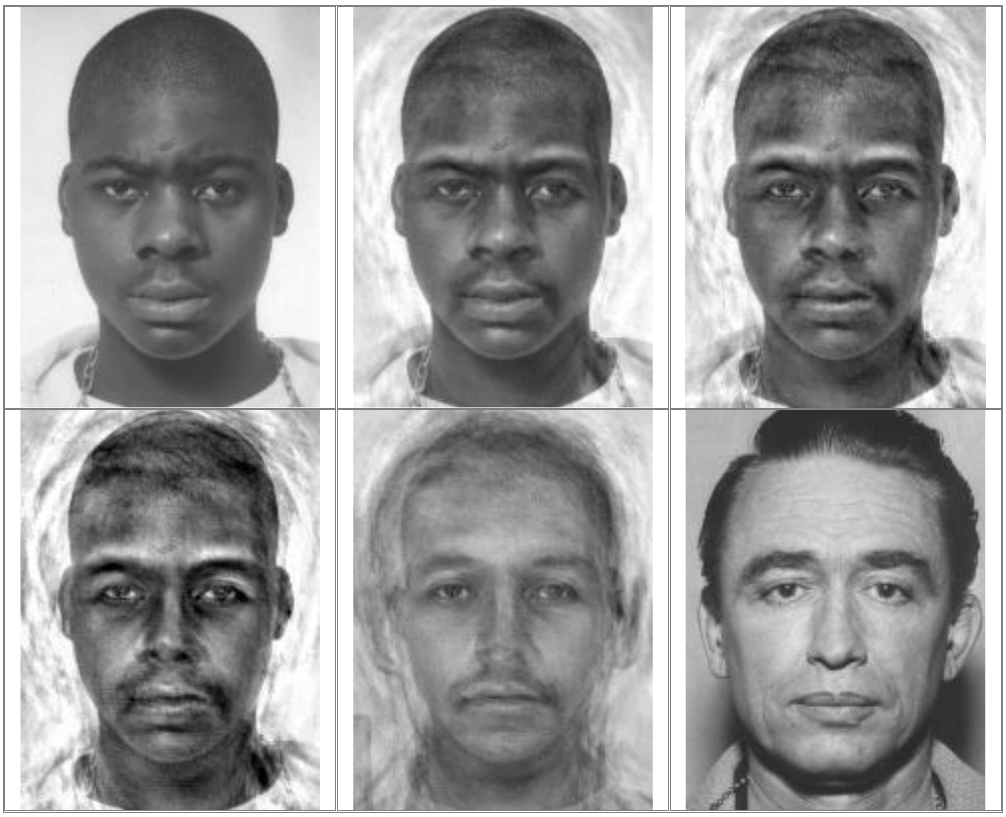

Figure 4: From left to right, top to bottom, estimated image, Im[k], at various iterations (k=0, 200, 500, 3200), average image from four different starting estimates, and the image of the target person.

4. DISCUSSION

An algorithm is presented which is able to substantially recreate the image of a person from the face recognition template. The only information required is the identity of the target template and the ability to calculate the match score between this template and an arbitrary image. The images calculated using the procedure are of sufficient quality to: 1) masquerade to the algorithm as the target, and 2) give a good visual impression of the person's characteristics.

This algorithm is not particularly sensitive to its details. Initial images that are visually distinct from the target converge to a good estimate; except for an increased calculation time. The minimization algorithm presented is simple and quite inefficient. There exist many more sophisticated multidimensional

minimization algorithms, which could presumably produce an image with less match score computations. The requirement for the algorithm is that it not be sensitive to local minima. The most important consideration is a good choice of basis functions. For face recognition, the eigenimage basis performs well, while basis functions with spatial discontinuities did not. In order to extend this technique to fingerprint recognition, where biometric features are spatially localized, a different type of basis functions would be required.

This technique could potentially be used in several real world scenarios. Recreation of a look-alike sample image is possible, and could be used to masquerade as the target or to identify them. For example, if an agency is allowed to compare their records against a biometric database of wanted persons, it would be possible for the agency to recalculate images of these persons. However, this approach does not allow exact recreation of the target image. This is not surprising, as the template typically contains significantly less data than the original image.

Recently, Hill [6] has shown that the templates from a fingerprint algorithm can be "hacked" in such a way as to allow synthetic images to be generated which would fool the algorithm to consider them the same as the original image. This work differs from the technique of Hill [6], in which the structure of a fingerprint template was expertly analysed. Firstly, access to the template storage is not required. This implies that encryption of the template [7], as proposed to counter the technique of Hill, will not work. Additionally, this algorithm does not require significant technical expertise to analyse a proprietary file format. Unlike a well designed encryption system, where the presented password is either correct or not, a biometric match score provides measure of closeness. It is this information which allows an initial guess to be gradually refined until it matches the target.

In conclusion, we have developed an algorithm by which an image of a person may be reconstructed from face recognition biometric match score results. The simplicity of this algorithm suggests that it be extensible to other biometric modalities. The implication in terms of privacy and security is that access to biometric results can effectively allow access to identifiable source images.

5. REFERENCES

[1] Adler, A., "Automatic Face Recognition System Architecture for Collaborative Virtual Environments", *IEEE Int. Workshop Haptic Virt. Environ.* 1:1-6, 2002.

[2] Arendt, H., "Biometric Identification and National Security", *Secure,* 5:56-57, 2002. (Internet: www.silicon-trust.com/ background/mag.htm)

[3] Argus Solutions, "Technology Overview – Privacy", *Internet:* www.argus-solutions.com/Privacy.html

[4] Ethentica Corp. "FAQ: Biometric Template Information", *Internet:* www.ethentica.com/template.html

[5] Grother, P., "Software Tools for an Eigenface Implementation", *National Institute of Standards and Technology,* Technical Report with FERET Database, 2000.

[6] Hill, C.J. "Risk of Masquerade Arising from the Storage of Biometrics", *B.S. Thesis, Austrailian National University,* 2001. (*Internet:* chris.fornax.net/biometrics.html)

[7] International Biometric Group, "Generating Images from Templates", *I.B.G. White Paper,* 2002. (*Internet:* www.ibgweb.com/ reports/ public/reports/templates_images.html

[8] Wayman, J.L. "Fundamentals of Biometric Authentication Technologies", *Proc. CardTech/Secure Tech,* 1999 (Internet: www.engr.sjsu.edu/biometrics/nbtccw.pdf)