

IMAGES CAN BE REGENERATED FROM QUANTIZED BIOMETRIC MATCH SCORE DATA

Andy Adler

*School of Information Technology and Engineering,
University of Ottawa
adler@site.uottawa.ca*

Abstract

We address the possibility of regenerating sample images from stored biometric data, specifically from automatic face recognition algorithms. Such algorithms calculate a match score from comparison of a newly acquired image of a person to a template calculated from previously captured images. Although several vendors of biometric algorithms claim that an image of a person cannot be regenerated from the template, it has been shown that, in general, such regeneration can be performed with a "hill climbing attack" (Soutar et al., 1999). In order to defend against this attack, the BioAPI consortium (2001), recommended that biometric algorithms emit only quantized match scores. In this paper, we show that it is still possible to regenerate biometric images even if the BioAPI recommendation is implemented. Each iteration of the algorithm is applied to a quadrant of the sample image. Before each calculation, noise is added to the image in the opposite quadrant, in order to force the match score to a value just below the quantization threshold. This means that the quantized match score is brought into a range where it provides useful information. Results show this algorithm successfully regenerates images which compare at high match scores for reasonable values of the quantization level. We conclude that the quantization of match score values does not, by itself, protect against the regeneration of images from stored biometric data.

Keywords: *Image Processing, Biometric identification, Face Recognition.*

1. INTRODUCTION

There is increasing interest in biometric authentication for government security applications. Technologies such as automatic face recognition, fingerprint and iris identification are being implemented for government identification documents and surveillance [13]. Such use of biometrics has provoked 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 [16]. 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. Many vendors of biometric systems have claimed that it is impossible or infeasible to recreate the image from the templates (for example [4,7]). In light of this claim, biometric templates may be considered non-identifiable data, much like a password hash [10]. Such biometric data can then be managed in ways that the source images cannot. For example, this assumed non-identifiability has been 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.

In fact, it is known that biometric templates are identifiable. Work by Soutar et al. [14] and ourselves [2] has shown that an estimate of the source image can be generated from biometric templates using a "hill-climbing" attack. This approach requires access to match score values for comparisons between an arbitrary image and the template of the target image. Beginning with a generic initial image, such an algorithm makes small modifications, and measures the resulting changes in the match score value. Modifications which increase the match score are retained, and, eventually, the modified image will resemble the unknown source image. In order to prevent such image regeneration, the authors of the BioAPI specification [5] recommend that match score values be quantized. Such quantization means that small changes to an image will normally not result in a change in the match score, preventing such a "hill-climbing" attack [14]. The level of quantization plays an important

role; a too-large spacing between quantization levels reduces the ability of application software to interpret the match-score data.

This paper develops a modified "hill-climbing" attack which is able to overcome such quantization of match score values. Results are shown for templates calculated with a face recognition algorithm.

2. METHODS

A software application was implemented with the goal of regenerating a face image of a target (IM_{targ}). Here, the *target* is defined as the person whose image is to be regenerated. The application has access to a local database of face images, and has network access to a biometric application server [1], in which the target template is stored. The software begins with only the identity of the target and the ability to obtain match scores (MS) of the target compared to an arbitrarily chosen image (IM). We represent this function as:

$$MS_i = \text{biometric_compare}(IM, IM_{\text{targ}})$$

In order to build such a software application, an attacker would require access to the target template, and the software to implement *biometric_compare*.

2.1 Hill-climbing algorithm

This section describes the hill-climbing algorithm for image regeneration using non-quantized match score data, based on the work of [2] and [14].

1. *Local database preparation*: A local database (LI) of frontal pose face images is obtained. Images 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.
2. *Eigenface calculation*: Using the local image database, LI , a principle components analysis decomposition is used to calculate a set of eigenimages (or eigenfaces [15]), using the method of [9]. The i th eigenimage is represented by EF_i .
3. *Initial image selection*: An initial estimate (IM_0) is chosen which is subsequently iteratively improved in the next step. A selection of images is chosen randomly from the local database, LI , and individually compared to the target. IM_0 is selected to be the one with the highest match score.
4. *Iterative estimate improvement*: Iterate the following steps for step number i .
 - A. Randomly select an eigenimage, EF_k
 - B. Iterate for step number j for a small range of values c_j , and calculate
$$MS_j = \text{biometric_compare}(IM_k + c_j \times EF_k, IM_{\text{targ}})$$
 - C. Select j_{max} as the value of j for which MS_j is maximized.

$$D. \text{ Calculate } IM_{i+1} = IM_i + c_{j_{\text{max}}} \times EF_k$$

E. Truncate values to image limits (ie. 0 to 255) if any pixel values of IM_{i+1} exceed the limits.

Repeat iterations until there is no significant improvement in match score.

The local database does not need to resemble the target image, and may be one of the many freely available face image databases [6,8,11,12].

2.2 Modified hill-climbing algorithm

The algorithm of the previous section does not work if *biometric_compare* returns quantized match scores [5]. This is because the small modifications to IM made in step B will not generally result in a change in MS . This section presents a modified algorithm which is able to function successfully with such quantized MS values.

1. *Local database preparation*: A local database (LI) is obtained and images are normalized as before.
 2. *Eigenface calculation*: An eigenimage decomposition of LI is calculated as before. The image is then equally divided into four quadrants (top left, top right, bottom left, and bottom right). Quadrant eigenimages ($EF_{i,\text{quadrant}}$) are then defined to be equal to EF_i within the quadrant and zero elsewhere. The edge of each quadrant is then smoothed to provide a gradual transition over 10 percent of the image width and height.
 3. *Initial image selection*: An initial estimate (IM_0) is chosen as before.
 4. *Iterative estimate improvement*: Iterate the following steps for step number i .
 - A. Randomly select an eigenimage, EF_k
 - B. Randomly select a quadrant Q . The diametrically opposite quadrant is referred to as OQ .
 - C. Generate an image RN consisting of random Gaussian noise in OQ and zero elsewhere.
 - D. Calculate the amount of contribution of RN which reduces the quantized match score by one quantization level.
 - Calculate $MS_i = \text{biometric_compare}(IM_k, IM_{\text{targ}})$
 - Using a bisection search, calculate the minimum value n to produce a noisy image, NI , where
$$NI = IM_k + n \times RN$$
such that the corresponding match score
$$MS_{NI} = \text{biometric_compare}(NI, IM_{\text{targ}})$$
is less than MS_i
 - E. Iterate for step number j for a small range of values c_j , using the quadrant Q eigenimage.
$$MS_j = \text{biometric_compare}(NI + c_j \times EF_{k,Q}, IM_{\text{targ}})$$
 - F. Select j_{max} as before.
 - G. Calculate $IM_{i+1} = IM_i + c_{j_{\text{max}}} \times EF_{k,Q}$
 - H. Truncate pixel values to image limits as before.
- Repeat iterations until there is no significant improvement in match score.

Values of c_j were selected heuristically for fastest convergence; the maximum value of c represented approximately 10% of the standard deviation of target image pixel values. There is a compromise in terms of convergence time between the number of iterations and number of values of c_j . This paper used 3000 iterations and 15 values of c_j (including zero).

This algorithm works separately on quadrants of the image. Because the quantized match score will not normally give information to allow hill climbing, a carefully chosen level of noise is introduced into the opposite image quadrant, in order to force the quantized score into a range where its information can once again be used.

3. RESULTS

Results in this paper were calculated using a commercially available face recognition software package. Initial tests with three other similar software systems show comparable results. Target and source images were chosen from the NIST Mugshot Database [12], and LI was calculated using the University of Aberdeen face recognition database [6].

Since the match score values are specific to any biometric algorithm, they do not show the significance of a result. In order to interpret MS values, results are shown in terms of the *confidence* of a genuine match. We define

the *confidence* to be the likelihood (in the Bayesian sense) that a comparison was genuine, given a match score MS was obtained. This statistic was estimated by performing comparisons between all possible genuine and impostor pairs of images in the database. For this algorithm, a quantization level of 1.0 corresponds to a maximum change in *confidence* of 10.6%.

The modified hill-climbing algorithm was applied to five different randomly selected source images at three different quantization levels: 1.0, 0.5 and 0.001. The latter level is effectively the same as no quantization. Each run of the algorithm required 135,000 biometric comparisons, and took 122 minutes on a 2.8 GHz Pentium IV PC computer. The 100 lowest order eigenimages were used for image regeneration. Figure 1 shows a graph of *confidence* versus iteration number for a representative image at each quantization level. In all cases, the algorithm is able to significantly improve the image in terms of its similarity to IM_{targ} . The initial *confidence* was 0.229. Without quantization, the algorithm is able to quite rapidly achieve high *confidence* matches (0.998). As the quantization level increases the algorithm slows and achieves a poorer best estimate (0.997 and 0.978 for quantization levels 0.5 and 1.0, respectively). However, even at the largest quantization level, the algorithm is able to produce an image which, in all cases, achieves a *confidence* above 95%.

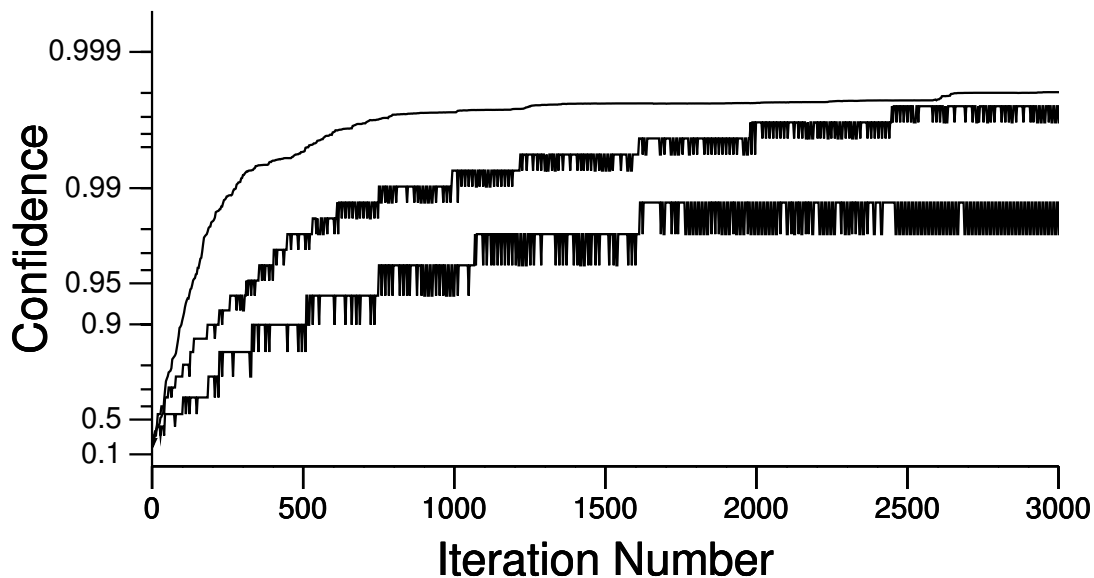


Figure 1: *Confidence* (calculated from match score value) versus iteration number for the modified hill climbing algorithm for a representative initial image (Top curve, quantization = 0.001, Middle Curve, quantization = 0.5, Bottom Curve, quantization = 1.0). Increasing quantization decreases the maximum *confidence* achieved by the algorithm.

4. DISCUSSION

This paper considers one aspect of the security and privacy implications of biometric data storage. Although some biometric algorithm vendors claim that an image of a person cannot be regenerated from a template, it has been shown that, in general, this can be accomplished using a "hill climbing attack" [2,14]. In order to defend against this attack, the BioAPI consortium [5], recommended that biometric algorithms emit only quantized match scores.

In this paper, we show that it is still possible to regenerate biometric images even if the BioAPI recommendation is implemented. The test image is split into four quadrants, and each iteration of the hill climbing attack is applied to one quadrant at a time. Before each calculation, the image in the opposite quadrant is "made worse", such that the match score is just below the threshold; after the calculation, the opposite quadrant is returned to its previous condition, and the best modification retained. Thus, the image is modified such that the quantized match score provides useful information. This algorithm was tested for three different quantization levels, and shown to be able to successfully regenerate an image which verifies at a high *confidence* to the original. As the quantization level increased, the algorithm showed a clear increase in run time and decrease in maximum *confidence* obtained. With a quantization level corresponding to a 10.6% change in *confidence*, the calculated image would not be able successfully masquerade against a system with a false accept rate setting of 1.0%. On the other hand, such a severe setting for the quantization level removes a significant amount of information from the calculated match score values. Furthermore, we anticipate that significant improvements are possible to this algorithm. First, only the lowest order 100 eigenimages were used, and the algorithm was stopped after 3000 iterations. We anticipate that increasing these limits will increase the maximum confidence at the expense of increased computation time. Secondly, the arbitrary division of the image into four quadrants, and the use of Gaussian random noise for degradation may not be optimal.

We conclude that the quantization of match score values does not protect against the regeneration of images from stored biometric templates. This result has privacy and security implications for the storage and transmission of biometric data, including biometric templates and match score values.

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*, 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 (current Mar. 2004)
- [5] The BioAPI Consortium, *BioAPI Specification (Version 1.1)* March, 2001. <http://www.bioapi.org/BIOAPI1.1.pdf>
- [6] 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.
- [7] Ethentica Corp., *FAQ: Biometric Template Information*, <http://www.ethentica.com/template.html> (current Mar. 2004)
- [8] 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.
- [9] P. Grother, "Software Tools for an Eigenface Implementation", *National Institute of Standards and Technology*, 2000. <http://www.nist.gov/humanid/feret/> (current Mar. 2004)
- [10] International Biometric Group, "Generating Images from Templates", *I.B.G. White Paper*, 2002. http://www.ibgweb.com/reports/public/reports/templates_images.html (current Mar. 2004)
- [11] 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 Mar. 2004)
- [12] NIST, *NIST Special Database 18: Mugshot Identification Database (MID)*, <http://www.nist.gov/srd/nistsd18.htm> (current Mar. 2004)
- [13] P.J. Phillips, "Human identification technical challenges", *Proc Int. Conf. Image Proc.*, 1:49-52, 2002.
- [14] C. Soutar, R. Gilroy, A. Stoianov, *Biometric System Performance and Security Conf.* IEEE Auto. Identification Advanced Technol., 1999. Also http://www.bioscrypt.com/assets/security_soutar.pdf
- [15] M.A. Turk, A.P. Pentland, "Eigenfaces for recognition", *J. Cognitive Neuroscience*, 3(1):71-86, 1991.
- [16] J.L. Wayman, "Fundamentals of Biometric Authentication Technologies", *Proc. Card Tech/Secure Tech*, 1999.