Can images be regenerated from biometric templates?

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

Increased use of biometrics has encouraged increasing concern about the privacy and security implications of these technologies. This paper considers 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. Enrolled biometric records are stored in the format of templates – a (typically vendor specific) compact digital representation of the essential features of the sample image. Because, biometric algorithm vendors have uniformly claimed that it is impossible or infeasible to recreate the image from the template, templates are traditionally considered to be *non-identifiable* data, much like a password hash. These claims are supported by the fact that: 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 difficult to "hack".

Recently, Hill (B.S. Thesis, Australian National University, 2001) demonstrated that is was possible to reverse engineer the file format of a particular (unspecified) fingerprint template. Software was developed to generate an image which would compare at high match score with the original, and would visually elucidate the essential characteristics of the original fingerprint. The implications of this work was analysed in a report by the International Biometric Group, where the following 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. This analysis recommended encryption of templates and the use of trusted devices to increase protection from image recreation from templates.

This paper presents an approach which expands on the on the idea of Hill. A simple algorithm is presented which can a regenerate sample images from templates using only match score results. While results are demonstrated for face recognition algorithms, the conceptual framework should be applicable to any biometric algorithm. 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 arbitrary face images, and has network access to a face recognition server. The software begins with only the database ID of the target person, and is able to obtain match scores of chosen images compared to the target person. Three different facial recognition algorithms were studied; all are recent products by well known commercial vendors of biometric systems. Two of the vendors participated in the face recognition vendor test 2002.

This algorithm functions as follows: During preprocessing, a local database of face images is obtained, and an eigenface decomposition calculated. Note that there is no requirement that the local database resemble the target image, these results use target images from the Mugshot face database, and local database calculated from the University of Aberdeen face database. Subsequently, the algorithm determines the match score for a selection of images in the local database against the target. The initial estimate is selected to be the image with the highest match score. The core of the algorithm is the iterative improvement of this estimate to better approximate the target. During each iteration, an eigenface image is selected, and a series of images produced equal to the current image estimate plus a small constant times the eigenimage. The corresponding match scores between these images and the target 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. 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.

Face recognition algorithms were tested for ten different initial images. The calculated match score showed a dramatic increase with iteration; in all cases, the result could successfully masquerade as the target image at a false accept rate of 0.1%. The figure shows representative initial images and the changes calculated by this approach for one face recognition algorithm.

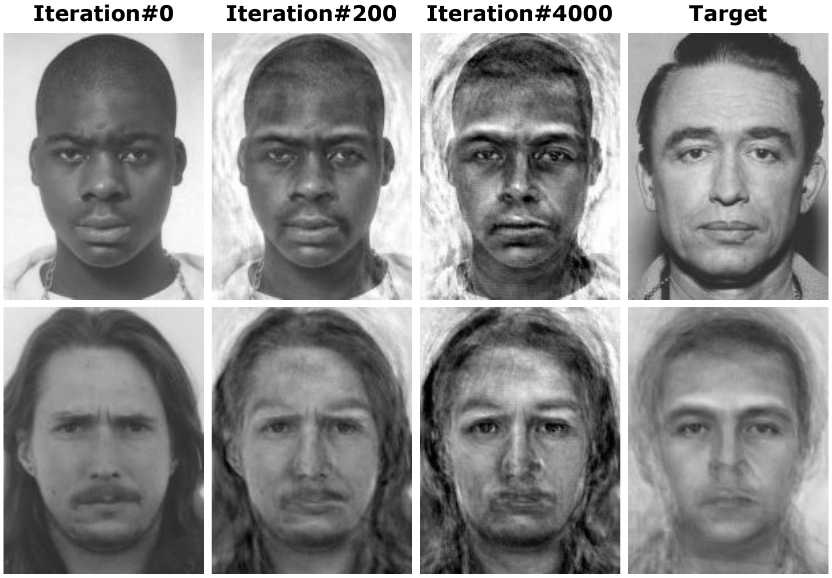

The target face (unseen by the algorithm) is on the top right. Two different initial estimates (iteration #0) are shown on the left side, next to the estimates at iterations 200, and 4000. The average estimate from ten different initial faces is shown on the lower right. Corrections to the image occur 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 average image (bottom right) has a surprisingly good resemblance to the target – even though it is quite blurred, and appears younger.

In conclusion, at least in the case of face recognition templates, a fairly high quality image of a person can be automatically regenerated from templates for three different algorithms, using only match score data. Importantly, this approach does not require a technically sophisticated user (as did Hill), and cannot be protected against by encryption of the template. 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. The simplicity of this algorithm suggests that it be extensible to other biometric modalities. This work implies that biometric templates and biometric match scores be considered identifiable data – they should not be made available to untrusted parties.