Robust distance measures for face-recognition supporting revocable biometric tokens.

T. Boult
University of Colorado at Colorado Springs and Securics, Inc
tboult@vast.uccs.edu

Abstract

This paper explores a form of robust distance measures for biometrics and presents experiments showing that, when applied per "class" they can dramatically improve the accuracy of face recognition. We "robustify" many distance measures included in the CSU face-recognition toolkit, and apply them to PCA, LDA and EBGM. The resulting performance puts each of these algorithms, for the FERET datasets tested, on par with commercial face recognition results.

Unlike passwords, biometric signatures cannot be changed or revoked. This paper shows how the robust distance measures introduce can be use for secure robust revocable biometrics. The technique produces what we call Biotopes™, which provide public-key cryptographic security, supports matching in encoded form, cannot be linked across different databases and are revocable. Biotopes support a robust distance measure computed on the encoded form that is proven to not decrease, and that may potentially increase, accuracy. The approach is demonstrated to improve performance beyond the already impressive gains from the robust distance measure.

1.0 Introduction and background

The paper has two major contributions, one on the use of robust distance measures for face-recognition, and the second on the issues of privacy. The two results are deeply interrelated, because as we attempted to solve the privacy issues, we realized the need for, and advantages of, a windowed robust operator in computing distance. This led us to the design, implementation and evaluation of robust measures for face recognition. The privacy issues are significant in their own right, though some may consider them outside of the scope of "vision". We contend, however, *our* community must address the many facets of our research, including societal impact. We briefly introduce some of the privacy issues before presenting, in section 2, the proposed approach. Section 3 presents more discussion of the application of robust distance measure in face recognition and then section 4 presents the experimental results.

A compromised biometric cannot be "replaced" and that permanent loss feeds the perception of invasion from any use of biometrics – if decades later the government or a corporation plays Big Brother, you cannot take back the information they gathered or lost.

On the security side, however, biometrics provide a unique way of verifying an individual and have very important uses. Proper biometric use can actually enhance privacy, by ensuring a verified identity before releasing data. A critical issue in the biometric area is the development of a technology that allies the privacy concerns while supporting the security goals.

This paper introduces the concept of biometric-based tokens that support unique identification, that support robust "similarity" or distance computations, that provide cryptographic security such that it can be canceled or revoked and replaced with a new one. This approach provides privacy while not compromising security. The critical idea is a technique that both provides a Public-Key-invertible mapping that hides the user's identity while simultaneously supporting a robust distance metric that allows the detailed matching needed to separate intra-subject variations from inter-subject variations and to support a range of different False Match Rate (FMR) vs. False NonMatch Rate (FNMR).

A few approaches for privacy preserving biometrics have been discussed in the literature, the most notable being [Ratha-et-al-01], wherein the biometric undergoes predefined distortion on the raw data (e.g. image) during both enrollment and verification. Such transforms can degrade the systems ability to detect the features needed for identification and can have a significant impact on the measure between the probe and gallery image. No accuracy impact on their approach has been presented. Some have presented "hashing" as a privacy protection, e.g. [Tulyakov-et-al –04], but to date those techniques almost doubled the error rates. In addition, the space of potential effective biometric consistent hashes appear to be small, resulting is limited privacy protection. Other related "privacy" work are the many papers/patents that improve privacy by mixing a user specific "random" pattern or phase mask that is mixed with or used to project the data, [Soutar-et-al-98, Savvides-et-al-04, Gao-Ngo-03]. These approaches, and others, apply a user-specific key to transform the data. [Goa-Ngo-03] does report the performance impact of their privacy improving techniques, but as their goal is "cryptographic key-storage", not recognition, it is not directly comparable. Furthermore, they only test on frames from a video stream that are very similar, not on any standard database. [Savvides-et-al-04] reports performance over illumination but the subset of the PIE data [Sim-et-al-01] used has no variation in pose/expression. In both papers it appears as a verification context, where the individual users key is used to generate their mask. It is unclear what the performance, or privacy protection, would be if the user key was known. In addition, [Soutar-et-al-98 and Savvides-et-al-04] provide correlation-like output that are all susceptible to hill-climbing attack, [Adler-05], to recover the original data. The strong

"windowing" of the approach proposed herein makes hill-climbing very difficult, as the distance landscape is constant except in a very small neighborhood of the true subject.

2.0 Secure robust biometric transforms

We will introduce the new approach; review the concept of robust distance measures and robust distance calculation. We will then briefly discuss issues of enrollment, transform storage, how to extend the approach to N dimensional data and issues for finite bit data. We then discuss a few variations on the encoding.

Our approach uses feature space transforms based on the representation of the biometric signature, i.e. after all transforms are computed. Most importantly the transform induces a robust distance/similarity metric for use in verification. In a sense, it is an "add-on" after all the other processing. The approach supports both transforms that are public-key cryptographically invertible, given the proper private key or using cryptographic one-way functions such as MD5 which trade less risk of comprise for more effort in reenrollment or transformation if data is compromised. In either case, even if both transformation parameters and transformed data are compromised, there is no way to get back the original data, thus removing the risk of reconstruction if centralized Databases are compromised. Finally, the approach can support an integrated multi-factor verification wherein the stored data cannot be used for identification (or search), even using the "guess each person and verify" approach. Existing multi-factor approaches store the biometric and other factors separately, verify each and only provide access if all are successful. Our approach stores a fused data and neither the biometric nor the added factors are directly stored in the DB. Thus allowing a face biometric for "verification" that the government cannot use for surveillance/search!

2.1 Robust distance computation

For the sake of simplicity in understanding, we initially explain the approach presuming all fields are 64 bit floating-point numbers. We illustrate the idea with a simple biometric signature with one field and we assume for simplicity of explanation that the "distance" measure is simply the distance from the probe to the gallery data (i.e. items in the DB) and that the "verification" is then based on threshold of the absolute distance.

A key insight into the approach is that a robust distance measure is, by definition, not strongly impacted by outliers [Huber-81]. In a robust measure, the penalty for an outlier is generally constant. Figure 1 shows the penalty function for a 1D least squares error and for 1D robust M-estimator. In many of the traditional distance measures, e.g. L2, weighted L2 or Mahalanobis measures, the multi-dimensions penalty for a mismatch grows as a function of distance, thus if the data in one

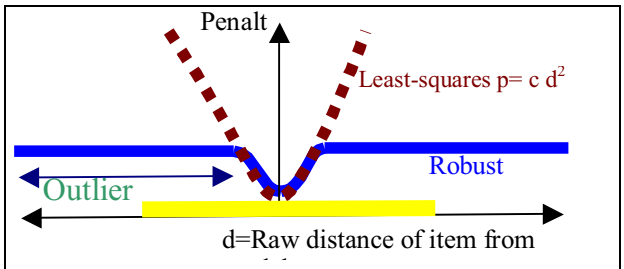

Figure 1: The penalty in the "similarity function". For weighted least squares errors, the penalty is a constant time distance, and grows quadratically. Thus a single outlier significantly impacts the fitting. For a robust similarity metric, the penalty is limited to maximum value e.g. outliers have a constant, and limited, impact on the overall measure. Given measurements p,q, we can define a robust measure $m_b(p,q) = c$ if $abs(r(p) - r(q)) > b$, and $m_b(p,q) = (r(p) - r(q))^2$ otherwise.

sub-dimension is far off then the penalty is high. Most fingerprint systems use a robust distance measure, yet most open-source face systems have only limited forms of them. (Most commercial face systems do not detail their similarity measures.) There are many other uses of the term "robust" in vision; here we stick to distance measures. Robust distance measures have been used in a range of vision problems many of which are parameter estimation (e.g. pose, stereo, motion), with nice range covered in special issue with introduction by [Meer-et-al-00] and a good review by [Stewart-99]. The approach herein is quite different from "Robust PCA", such as [Huber-et-al-05]. We do not modify the PCA computation, only the distance measure use for "classification" in the PCA/LDA or other algorithm subspaces.

There are many types of robust distance. In the remainder we consider only windowed operators, where the penalty outside a finite window is constant (or 0). Before we look at privacy preserving, we explain how this type of distance measure can enhance a recognition algorithm. Consider the hypothetical clusters show in Figure 2, with the axes showing the first two PCA coefficients. Recall Mahalanobis space is "defined" as a space where the sample variance along each dimension is one, which is the case in Figure 2. It is important to note that the overall data co-variances are quite different from the per-cluster co-variances. The text labels correspond to the ideal clusters of a collection of sample instances. "Ideal" elliptical boundaries are shown for each cluster. The background circle with wavy lines represents a constant distance from the center of the "1" cluster. Using simple distance, the tilted text shows the incorrect associations from clusters 2 and 3 with the "1' cluster, as well as the incorrect associations of two 1's with the 2 and 3 clusters.

To consider the robust-distance measure's impact, let dashed lines represent the diameter of each dimension of a "robust window", implicitly rescaling distances so that

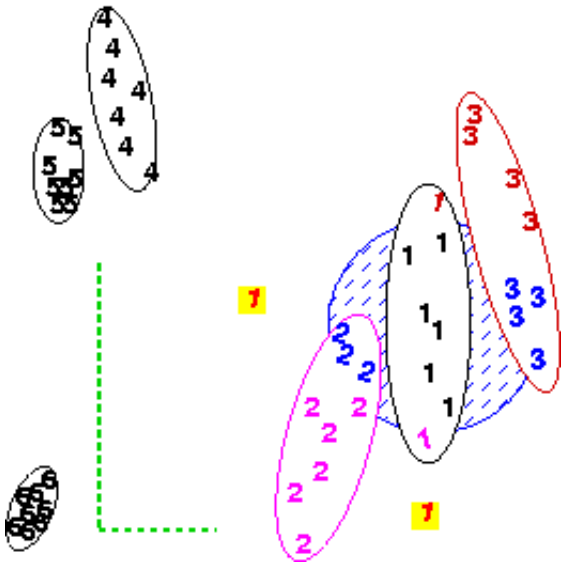

Figure 2: Cluster association using traditional and robust distance measures. See text for discussion.

this is a unit length in each dimension. This window size is the maximum extent of the non-outliers (i.e. excluding text in boxes) for any cluster. In this example, clustering using the robust distance measure provides a perfect separation of the clusters. We note that even the hypothetical "outliers" (text in shaded boxes), properly associate using this robust distance operator. Outliers are especially problematic for projection-based approaches, as a small alignment error can cause significant error in projections with high-frequency patterns in the subspaces coefficients.

Note this example considered a single "robust window" size for all data, a process we call GroupRobust. Further clustering improvements can be realized, if we have sufficient training data, by using a separate robust window per cluster. In either case, this robustification is a highly nonlinear process and requires appropriate training. The size of the robust-window should be the expected size of intraclass variations, ignoring outliers. Of course, we cannot presume such powerful a priori knowledge, so an important aspect of the approach is how well experimental data allows us to estimate the needed window sizes.

We now return to the issue of privacy. We assume the biometric produces a value v which we then transform via scaling and translation, v'=(v-t)*s. The resulting data is separated into two parts, one representing overall position and the other local detail or residual. Without loss of generality, we can represent this with residual r = the fractional part of v', and general wrapping number g = the integer part of v'. The approach will then use a one-way transform or encryption of g to hide the user's identity. The thick shaded line on the axis of Figure 1 shows an example "residual region" after appropriate transforms. After the transform, all data will be "wrapped" or aliased to some position r within this region. The number of times the data is wrapped is the value g. This describes the key concepts of the approach. A mapping hides the actual value but as it separates the result into two, it leaves an unencrypted value within a "window" in which we can compute local distance, and then encrypts the larger (and hence very stable) part of the position information thus effectively hiding the original position and protecting privacy. This key idea was inspired from the modulus computation in RSA-type Public Key (PK) algorithms.

An alternative enhancement includes a user password before encoding. The transform and wrapping is computed and the passcode is then fused with the generalized wrapping index, g before encoding. The inclusion of the passcode provides a strong form of revocation, and protection from its use in search or identification rather than recognition.

To ensure that the biometric data is protected even if the "transformation" parameters are compromised, we need to ensure that the mapping from g to w is non-invertible or at least cryptographically secure. The "security" of the revocable approach is determined by this transform. The preferred approach is to use a PK encryption of g to produce w. For simplicity we refer to the transformation v to (r,w) as encoding and r,w as encoded data. If a password is mixed with g before encoding, the result is a revocable biometric that is suitable for verification but which prevents recognition or search as the user pin is not stored anywhere.

2.2 Robust distance computation on the encoded data

Assume for signatures p,q, encoding using s,t yielding r(p), r(q), w(p), w(q), we define the robust dissimilarity metric d(p,q) as follows:

$$d(p,q) = c \text{ if } w(p) \mathrel{!}= w(q) \;||\; abs(r(p) - r(q)) >= b$$

$$d(p,q) = (r(p)/s(p) - r(q)/s(q))^2 \text{ otherwise}$$

This distance computation is just one example of a robust distance measure, one that uses a constant penalty outside a fixed window and least squares penalty within the window. The unique property of the mapping ensures that the window around the correct data is mapped to a window in which any robust distance measure can be computed.

Clearly given r,s,t and g, the original data can be reconstructed. It should also be obvious that many distinct data points will all have the same value for r, and that without knowledge of g, the original cannot be recovered. The biometric store would maintain r,s,t and w (the encrypted version of g). We can consider each of these as user specific functions that can be applied to an input signature, e.g., $r_k(v)$ is the residual associated with biometric signature v when using the kth user's transform, and $w_k(v)$ is key w that results from v after applying the transform and the encryption associated with user k.

A key issue is the choice of the scale and translation. If we let $e_{k,j}$ be the jth biometric signature for user k, then if s_k and t_k are chosen such that

$$b\, s_k < r_k(e_{k,j}) < (1-b\, s_k) \quad \forall\, j \qquad \text{(Eq 1)}$$

for each field in the signature. Since we are free to choose s and t separately for each user and each field and can do so after we have obtained the enrollment data, it is straightforward to satisfy Equation 1 for all enrollment data. For this to be truly effective, the range of values used to determine the scale in Equation 1 should larger than the actual variations of that parameter for that user, not just over the enrollment data. In practice, we have increased the enrollment range by a factor of 3 to ensure that the actual user's data is very unlikely to fall outside the scaled window. Even with the described constraints there are still "infinitely" many choices for t for "real" numbers, and a huge range for floating points. Changing t impacts both r and g and combined with the encryption for w, provides protection of the underlying identity. For finite bit representations, the constraints are more limiting, as is discussed later, but for some values of b it can be satisfied for any field with more than a single bit.

Theorem 1:
If a transform satisfies Equation 1, and the distance measure is has a constant penalty outside the window that is at least as large as any penalty within the window, then computing distances in the encoded space cannot decrease, but may increase, the accuracy of the system.

Proof: Given Equation 1, it is easy to see that $d(p, e_k) = m_{sb}(p,e_k)$, i.e. for the matching users the robust dissimilarity measure applied to the transformed data is the same as the original robust metric applied to the raw data with a robust window of size (s_k* b).

For an imposter, e_i, encoded with user k's transform two possibilities exist. If $b\, s_k < r_k(e_{i,j}) < (1-b\, s_k)\ \forall\, j$, for every field within the signature then w(p) == w(q) and the distances for the imposter i, will be the same before and after transform. Otherwise, scaling/shifting has resulted in at least one field distance being equal to c, even though the field was initially close enough that the pre-encoded distance was < c. Since c is chosen such that it is greater or equal to the maximum distance within the robust window, then for non-matching i != k, the transform may increase, but cannot decrease, the distance. *Q.E.D.*

Thus we have shown the distance transform computed in encoded space cannot decrease accuracy, and may often improve accuracy by increasing the distance to non-matching subjects. Note the proof requires the scaling/translations satisfy Equation 1, which need not be the case, especially if are outliers. Note that some "robust distance" operators have zero impact (penalty) sufficiently far from the data, which means they do not satisfy the preconditions of the theorem.

For a real biometric with N dimensions, we treat each dimension separately, so given a raw biometric vector V with n elements, we compute V'=(V-T)*Diag(S), where each of T and S are now vectors of size N and Diag(S) is an N by N diagonal matrix generated from S. We separate the result of the transformation, this time into the residual vectors R, and general wrapping G. Again G is transformed to the encrypted W, and the biometric store retains T, S, R and W. If the system designer usually uses a Mahalanobis transform before distance computation, the covariance transform should be applied to V before it is subject to translation and scaling. Note this simplifies the process since after the covariance transform, the data is mean-zero and scaled so that the variance in each direction is equal. With such a transform it may be possible to choose a single scale parameter S rather than an independent scale for each field, thus reducing the extra storage requirements.

3.0 Robust distances for face-based biometrics

While the secure robust revocable biometric transform applies to almost any biometric template, we evaluate its performance on face-based systems. To demonstrate the generality of the improvements to be obtained by using robust distance measures in face-based systems, we extended algorithms included in the Colorado State University (CSU) Face Identification Evaluation System (Version 5.0) [Bolme-etal-03]. In particular we developed robust versions of the "baseline" PCA-based face recognition system using multiple metrics, their LDA-based face recognition algorithms and the Elastic Bunch Graph Matcher (EBGM)[Okada-et-al-98]. Unless noted otherwise, we use the defaults, e.g. 300 coefficients for PCA and 427 for LDA

By design, the proposed approach maintains the robust distance measure after encoding. Thus the first step needed was to develop a face-based recognition system that used a robust distance measure. A robust distance measure is one where the penalty for outliers is bounded, with many different robust measures used in practice [Huber-81].

Robust distance measures have been used in various other computer vision problems; we were surprised to find no formal application in face recognition. The most "robust" distance measure we have seen in face recognition is the Mahalanobis Cosine" (hereafter MahCos) used in CSU. This is robust in a formal sense since the impact of an outlier is bounded, because the cosine computation itself has a maximum value of 1. We note that the MahCos was the best performing distance measure in the CSU tests. As will be seen in section 4, the secure revocable transform applied to MahCos, did produce a measurable performance improvement, but for unknown reasons, its overall

performance was well below the other robust approaches.

While there are many different robust measures, we consider only simple windowed or trimmed measures introduced earlier. Let ω be a window operator such that $\omega(\beta)= \beta$ if $\beta < \tau$, and $\omega = C$ otherwise, where τ is the window size. For the PCA algorithm, we implemented robust versions of "Euclidean" or L2 metric, as well as the robust Mahalanobis L2 measure (hereafter MahL2) and the CSU MahCos. We also extended CSU's LDA implementation to use a robust distance measure. For brevity, we describe only MahL2.

Mahalanobis space is "defined" as a space with unit sample variance along each dimension. The Mahalanobis Transform thus divides each coefficient in the vector by its corresponding standard deviation. Let u and v represent the unscaled PCA coefficient vector of images of user j and k respectively. Let σ_i be the standard deviation across all users in the i^{th} dimension, then the Mahalanobis transforms of u and v are given as $m_i= u_i/ \sigma_i$ and $n_i= v_i/ \sigma_i$ n_i. Let $\beta_i = (m_i-n_i)^2$, and let $s_{i,k}$ be the enrollment scaling from Equation 1 for user k, we define RobustMahL2 distance as

$$D_{RM}(j,k) = \sqrt{\sum_i \omega(s_{i,k}\beta_i)} \quad (2)$$

Obviously because of the scaling this is not generally symmetric, i.e. $D_{RM}(j,k) != D_{RM}(k,j)$.

In any windowed robust measure, the choice of the window size is very important. In our case, to ensure the revocable transform has the same distance before and after encoding, we desire the enrollment process to determine a scale such that Equation 1 is satisfied.

While per individual scaling produces a better robust measure for that individual, it can be problematic in that it presumes a wide range of images per individual for enrollment. We postulated that for each field, a single scaling could be used for the entire population. This simplifies enrollment, allowing for single image enrollment, but does slightly reduce the effectiveness of the robust distance transform. We call this the GroupRobust transform. This approach has worked well for both PCA and LDA with different robust measures.

The EBGM was extended to also support robust distance measures. The CSU system had a plethora of distance measures applied in EBGM matching which we extended. We report only two CSU considered the best.

4.0 Evaluation using face-based biometrics

This section summarizes our experimental valuation of the secure robust revocable biometric transformation (SRRBT) and shows the significant advantages of using a robust distance measure for face recognition. As noted earlier, previous work on transformed biometrics did not provide quantitative evaluation of the recognition performance so we cannot compare to such transforms.

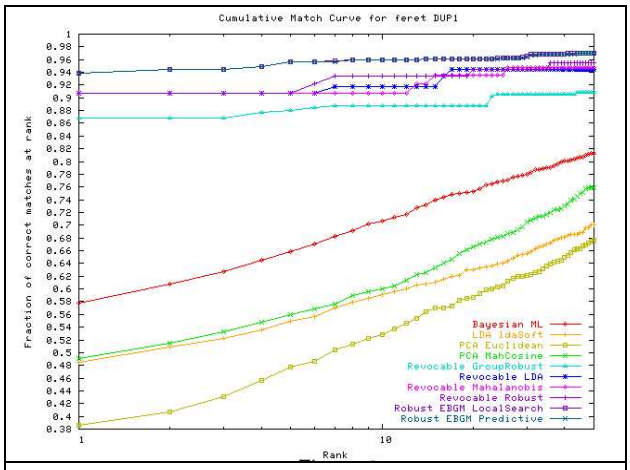

Figure 3: CMC curve for various algorithms on DUP_1

A robust revocable transform is first determined in section 2 for each individual and entered in the DB. For verification, the SRRBT of the claimed identity is applied to the probe data and then compared, using the robust distance measure, with the stored data. We treat identification/recognition as a sequence of verification attempts, apply person k's transform and then compute the distance to person k's revocable template.

In keeping with the CSU toolkit model, the experiment applied the robust revocable biometric to a gallery of all the FERET data to generate all pair-wise comparisons, and then subsets of that data were analyzed for different "experiments". The standard FERET experiments were done including FAFB, FAFC, DUP1, and DUP2 [Phillips-et-al-00]. The Secured Robust Revocable Biometric consistently outperformed the CSU baseline algorithms as well as all algorithms in the FERET study and all commercial algorithms with published results on FERET.

A CMC graph for some of the tested algorithms is shown in figure 3. This example CMC graph is similar to those for the other datasets. Hereafter we report only rank-1 data to save space. Table 1 shows the Rank 1 recognition rates computed for the standard FERET subsets for the algorithms in the CSU toolkit (gold), the best previously reported [NIST-01] from FERET tests (red) and a range of revocable robust techniques, with a total of over 250 million biometric comparisons. The dramatic improvement in recognition performance shows the significance of using a robust metric. To quantify the impact of the embedding separate from the robust distance measure, the tables include a "simple robust" which is the robust PCA algorithm without any embedding. An obvious issue for the GroupRobust techniques is the definition of the group used for training. We have tested with different groups, all 3541 images, DUP1 (243 subjects, 722 total images), FAFC images (2 each of 194), and the 2 images each of 71 individuals (X2) used to train the FERET PCA space (feret_training_x2.srt from CSU's toolkit), as well as

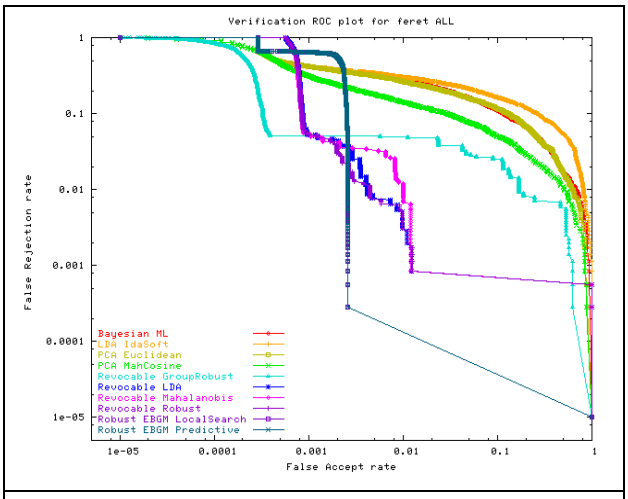

Figure 4: Verification ROC plots for FERET_ALL.

other subsets not shown. Note that FAFC has no subjects/images in common with any of DUP1, DUP2 or X2. Also note that differences of 1-2 individual recognitions (e.g. 100 versus 99.48 for FAFC) may be caused by the random "offsets" used to define the revocable transforms and should not be considered significant.

Verification rates are best computed over large sets, so in addition to the individual sets reported in the FERET study, we computed results for the full FERET_ALL, which has 3541 tests for positive identification and over 11 million attempts for false identification. The ROC in Figure 4 is in log-log format because the new algorithms perform so well that a linear ROC plot is useless. The graph truncates to 1e-5. Note for the revocable algorithms, the vertical axis of the ROC curves is generally truncated by the sample set size causing a radical jump to the right axis when the data is exhausted. The Equal Error Rate (EER) of the Revocable Robust algorithms that use individual scaling is generally better than prior algorithms by a factor of at least 10, often by 100. The Revocable Group Robust algorithm has a better equal error rate than all previously reported algorithms.

We also did experiments to show the revocable transforms do not allow matching/linking across databases or cannot be used without their associated passcode. The first self-matching test took 191 subjects in FAFC, and made 25 copies of each image, with the resulting rank one recognition being 0.0055 or about the 1 in 200 expected from random matching. We also processed the FERET_ALL set (3541 images) and after enrollment gave each image of a person a different transform. The resulting rank one recognition was zero on 3 out of 5 rounds and 0.0002 on the remaining 2 runs. For verification, the ERR for both tests was consistently 0.9997. Thus, as expected, different transforms for an individual match at random and hence protect privacy. It is important to point out, that in verification, the tests all provide the "imposter" with the correct "key" for generation. If the correct keys are only provided to the true match, the verification rate is driven by the random change of two keys colliding times the verification error rates presented above.

Algorithm	DUP1	DUP2	FAFB	FAFC
# subjects	243	75	1195	194
# Matched scores	479	159	1195	194
# Non-matched	228 K	25 K	1427K	37 K
PCA L2	33.79	14.10	74.31	04.64
PCA MahCos	44.32	21.80	85.27	65.46
LDA ldaSoft	44.18	18.80	70.96	41.75
EBGM Predictive	43.63	24.78	86.94	35.57
EBGM Search	46.26	24.35	89.79	41.75
FERET "BEST"	59.1	52.1	86.2	82.1
Simple Robust PCA	85.73	85.47	98.32	100.0
Revocable Robust PCA	90.72	87.18	99.50	100.0
Revocable (all) GroupRobust PCA	86.57	85.47	98.32	100.0
Revocable (DUP1) GroupRobust PCA	85.46	85.47	98.24	100.0
Revocable (X2) GroupRobust PCA	83.80	83.76	97.99	99.48
Revocable (FAFC) GroupRobust PCA	81.85	82.05	97.15	99.48
Revocable Robust PCA MahL2	90.72	87.18	99.50	100.0
Revocable Robust PCA MahCosine	68.14	67.52	93.97	96.39
Revocable Robust LDA	90.72	87.18	99.50	100.0
Revocable (all) GroupRobust LDA	88.78	85.47	98.91	100.0
Revocable (X2) GroupRobust LDA	87.95	84.62	98.83	100.0
Revocable (FAFC) GroupRobust LDA	81.85	81.20	98.24	99.48
Revocable Robust EBGM Predictive	91.27	88.03	100.0	100.0
Revocable Robust EBGM Search	91.27	88.03	100.0	100.0

Table 1: Rank 1 Recognition Rates on FERET subsets

5.0 Conclusions

The paper introduced the use of a windowed robust distance measure for face biometrics and showed how these measures significantly improve the performance of three well-known face-based biometric algorithms. Multiple different robust measures were tested per algorithm and all significantly improved performance. The "trimmed" robust measures used are different from

most used in previous vision because their penalty (loss) functions do not go to zero but to a large constant, and because of how the roust window sizes are estimated. However, if the training data is representative, their influence function does go to zero and the breakdown point is the ideal 0.5. The robust measures introduced require minimal added computational cost. The performance enhancements from the use of our robust transform were dramatic. The resulting Revocable Robust PCA, using multiple images for enrollment, produced performance on par with the commercial algorithms. The "group robust" versions, which permit use with only a single enrollment image, were slightly weaker but still better than most known algorithms and were good enough to produce 100% recognition on FERET FAFC when the group parameters are obtained from FERET DUP1 (a non-overlapping set of people). The excellent performance of the group robust algorithm based on PCA combined with the simplified enrollment processes and fast computation, combine to suggest this algorithm has considerable potential.

The new approach estimates consistent class "cluster sizes" and showed that even using different people for training, defining a robust windowed distance can provide significantly improved performance. A hypothetical explanation is that within an individual's "cluster", the variations represent non-linear interactions with unmodeled variable (such as pose, lighting and expressions). The subspaces for recognition consider the overall spread of people and unmodeled variables; The robust window distance separates out the latter. The GroupRobust results suggest the ranges of variations are relatively consistent across a population. We note "robustness" herein is quite different than that traditionally considered, e.g. as use in ROBPCA [Huber-et-al-05]. However, a combination of the two different approaches could be very interesting.

This paper introduced the revocable robust biometric transform and showed its effectiveness on face-based biometrics. The transforms are applied to biometric template data to produce two components one of which is encrypted while the other is stored unsecured. The transforms combined with encryption maintain the privacy while the unencrypted part supports a robust distance measure, something that is critical to make biometrics effective. *The paper proves that this privacy preserving transform will not decrease, and may increase, system accuracy.* While the paper presents only face, the approach applies to almost all biometrics and we have also implemented a fingerprint-based version. The fingerprint version is more complex because of the alignment issue, but still showed an average 30% improved accuracy.

Biometrics have the promise to improve security. But as Admiral James Loy, then Head of Transportation Security Agency, stated at the 9th Annual Privacy & American Business Conference, 2003 "Don't be too quick to strike a balance between privacy and security. As Americans, we are entitled to a full measure of both". Secure Robust Revocable Biometrics show, that at least for biometrics, we don't have to accept the loss of privacy to gain security. Not only do they provide privacy, they actually improve the accuracy of the underlying biometrics, which improves both privacy and security!

6.0 References

[Adler-05] Andy Adler, "Vulnerabilities in biometric encryption systems," Audio- and Video-based Biometric Person Authentication - AVBPA 2005, July 20-22. LNCS 3546:1100, 2005.

[Bolme-etal-03] D.S. Bolme, J.R. Beveridge, M. Teixeira, and B.A. Draper. The CSU Face Identification Eval. System: Its Purpose, Features, and Structure. *ICVS 2003*: 304-313.

[Goa-Ngo-03] A. Goh, D.C.L. Ngo, "Computation of cryptographic keys from face biometrics". Int. Fed. Info. Proc., Springer-Verlag, LNCS2828, pp. 1-13, 2003.

[Huber-81] P.J. Huber, *Robust Statistics*, John Wiley & Sons, New York. 1981

[Huber-et-al-05] Hubert, M., Rousseeuw, P.J., Vanden Branden, K. (2005), ROBPCA: a new approach to robust principal component analysis, Technometrics, 47, 64-79

[Meer-et-al-00] P. Meer, C.V. Stewart, and D.E. Tyler, "Robust computer vision: an interdisciplinary challenge", *CVIP* 78:1, April 2000. Introduction to the special issue on robust statistical techniques in image understanding.

[NIST01] www.itl.nist.gov/iad/humanid/feret/perf/score_cms/score_cms.html. Last downloaded 5/2005

[Okada-et-al-98] K.Okada, J. Steffens, T. Maurer, H. Hong, H. Neven, and C. von der Malsburg. The Bochum/USC Face Recognition System and How It Fared in the FERET Phase III Test. In *Wechsler et al.*, editors, *Face Recognition: From Theory to Applic.*, pp. 186-205. 1998

[Phillips-et-al-00] P. J. Phillips, H. Moon, S. A. Rizvi, and P. J. Rauss. The FERET evaluation methodology for face-recognition algorithms. IEEE TPAMI, 22(10):1090 1104, 2000.

[Ratha-et-al-01] *N. Ratha, J. Connell, R. Bolle "*Enhancing security and privacy in biometrics-based authentication systems**",** *IBM Systems Journal*, vol. 40, no. 3, pp. 614-634, 2001

[Savvides-et-al-04] M. Savvides, B.V.K. Vijaya Kumar and P. Khosla, "Authentication Invariant Cancelable Correlation Filters for Illumination Tolerant Face Recognition" Proc. of SPIE, Biometric Technology for Human Identification, Vol. 5404, Orlando , FL, 2004

[Soutar-et-al-99] C. Soutar, D. Roberge, A. Stoianov, R. Gilroy and B.V.K. V. Kumar, "Biometric Encryption," ICSA Guide to Cryptography, McGraw-Hill, 1999.

[Stewart-99] C.V. Stewart, "Robust Parameter Estimation in Computer Vision", SIAM Review 41:3, 1999.

[Tulyakov-et-al –04] S. Tulyakov, V. Chavan and V. Govindaraju. "Symmetric Hash Functions for Fingerprint Minutiae". Biometrics Consortium Conference, Crystal City, VA, 2004.

[Uludag-et-al-04] U. Uludag, S. Pankanti, S. Prabhakar and A.K. Jain,"Biometric Cryptosystems: Issues and Challenges", Proceedings of the IEEE, Vol 92, No 6, June 2004.