FAAD: Face At A Distance

Terrance. E. Boult[12], Walter J. Scheirer[1,2] and R. Woodworth[2],
[1]University of Colorado at Colorado Springs and [2]Securics, Inc

ABSTRACT

The issues of applying facial recognition at significant distances are non-trivial and often subtle. This paper summarizes 7 years of effort on Face at a distance, which for us is far more than a fad. Our effort started under the DARPA Human Identification at a Distance (HID) program. Of all the programmers under HID, only a few of the efforts demonstrated face recognition at greater than 25ft and only one, lead by Dr. Boult, studied face recognition at distances greater than 50 meters. Two issues were explicitly studied. The first was atmospherics/weather, which can have a measurable impact at these distances. The second area was sensor issues including resolution, field-of-view and dynamic range. This paper starts with a discussion and some of results in sensors related issues including resolution, FOV, dynamic range and lighting normalization. It then discusses the "Photohead" technique developed to analyze the impact of weather/imaging and atmospherics at medium distances. The paper presents experimental results showing the limitations of existing systems at significant distance and under non-ideal weather conditions and presents some reasons for the weak performance. It ends with a discussion of our FASST™ (failure prediction from similarity surface theory) and RandomEyes™ approaches, combined into the FIINDER™ system and how they improved FAAD.

1. SENSOR ISSUES FOR FAAD

While many people discuss face at a distance and use traditional NTSC type cameras, an actual application has to decide how the images of the subject of interest will be acquired. If it is a cooperative subject, then many assumptions can be made. But for a non-cooperative subject, the real focus of FAAD, the scenario must be focused on a choke point e.g. a door or pathway where the subject can be expected to appear. The system designer must consider the field-of-view, resolution and lighting issues when choosing components. A summary of effective FOV for different sensor resolutions is shown in Figure 1 and Table 1. For low light scenarios some researchers have discussed using thermal sensors [12], but we consider their resolution generally too small to even consider. Even standard cameras, as 640x480, provide a FOV too small to seriously consider for acquiring non-cooperative subjects at a chokepoint. FAAD systems designers must consider the resolution choice early as it impacts processing requirements, lens choices and lighting requirements.

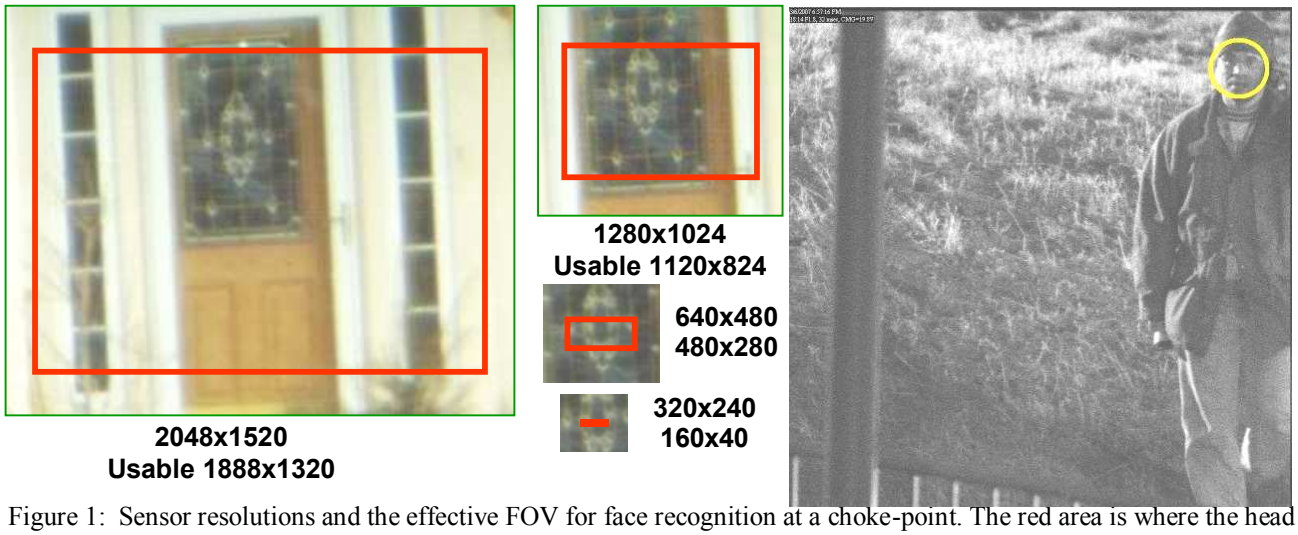

Figure 1: Sensor resolutions and the effective FOV for face recognition at a choke-point. The red area is where the head must be for the face to be usable assuming an optimistic 80 pixel "face" (50 pixel IPD). Remember the "usable" (red) area must also account for variations in subject height and thus a mega-pixel sensor is probably a minimum for usable resolutions. The image on the far right is an EMCCD image from 100m using a 400mm f2.8 lens on a moonless night with a streetlamp 50m on the left. The circle is the automatically detected face.

A critical sensor issue is that of lighting and sensor sensitivity. In applications where 24/7 operation is required, even with added illumination, some type of image intensification is needed. A study of traditional intensifier impacts, ignoring the FOV issues, can be found in [1]. In our own work we also tested intensifiers but found the spatial resolution loss from the intensifiers MTF, combined with its small FOV it provided was too limiting.

Sensor resolution	**2048x1520**	**1280x1024**	**640x480**	**320x240**
Usable size in pixels	1888x1320	1120x824	480x280	160x40
Choke point size (in ft)	5.85' x 4.3'	3.6' x 2.6'	1.5' x 0.9'	0.5' x .1'
Allowed height variation	± 25in	± 15.6 in	± 5.4 in	± 0.6 in
Diagonal size (in pixels)	2303	1390	555	164
Diagonal crossing time	2.8sec	1.2sec	.69sec	.2 sec
Vertical/forward crossing	1.65sec	1.03sec	.35sec	.05 sec

Table 1: Usable Resolution, size and time within FOV when FOV is maximum size that optimistically can be used for face recognition. Conservative estimates would be half the sizes/times shown.

For the past 2 years Securics and UCCS have been jointly developing our FIINDER™ (FPGA-enhaed Image Intensified Networked Detector with Embedded recognition) system. For that system we have switched to using an EMCCD camera [2], which we found to be superior to tube-based intensifiers. Even with an improved sensor there is a significant need for lighting normalization, especially to address the issues of the directional illumination. The low lighting and noise of intensified imagery, coupled with the potential for strong directional lighting (after amplification) makes this a challenge. While there has been some good work on directional lightning in the past, e.g. [12], we found that approach expensive and still very limited. Thus we developed a new (patent pending) lighting normalization algorithm just to address strong directional lighting artifacts and do so in a manner well suited to hardware implementation. Figure 3 shows an example of this applied to multiple low-light examples as well as comparing it to the standard normalization from the CSU face toolkit [3] and the approach from [12]. The basic concept of this "DUAL LUT" normalization is to do histogram shaping, a generalization of histogram equalization, separately around each "eye" to determine a desired normalization for that region, which is then implemented as a look up table (LUT). Then, each column of the (potentially rotated) image is normalized using a blending of the two normalizations with pixels between the eyes doing a bilinear blending of the two normalizations. Beyond the eyes, i.e. the outside of the face, the normalization from the closest eye is used. With two lookup tables and a spatial defined bi-linear blending, it is very well suited to a FPGA implementation. The quantitative impact of this normalization, and comparison with standard

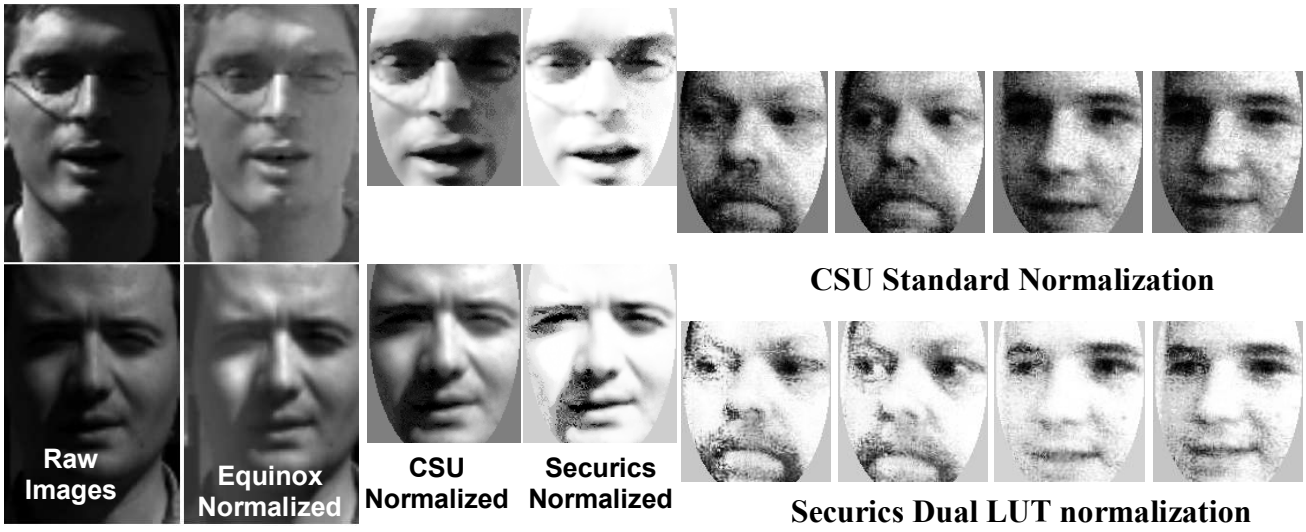

Figure 2: Examples of low-light imagery normalized. Left is an example image from [12] . The second column shows the Equinox normalization from that paper, the third column showing the standard CSU normalization applied to the images and the fourth column showing the Securics dual-LUT normalization. The right 4 columns show more examples comparisons of CSU normalization and Securics dual LUT normalization.

normalization, is discussed in section 5.

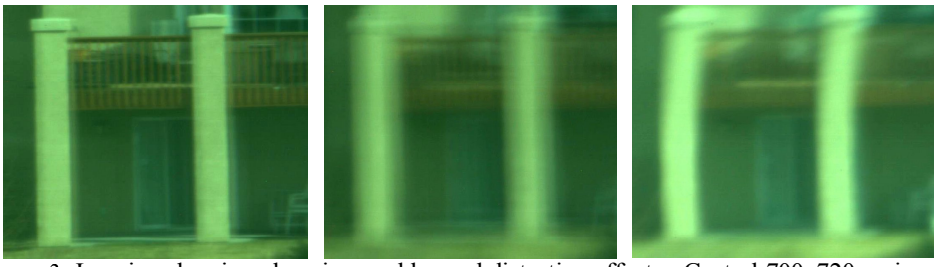

Figure 3: Imaging showing clean image, blur and distortion effects. Central 700x720 region from 2048x1520 image at 800m F8

Another important issue is the choice between rolling and frame shuttering, which in turn depends on the lens, and stabilization. A rolling shutter, even at moderate shutter speeds, allows for the motion within the frame to become a geometric distortion rather than just a blur. While blur is a problem, geometric distortions have an even more dramatic impact on FAAD. Figure 3 shows results, with a camera on a tripod, showing they types of geometric distortions that can occur. Some interesting results addressing de-blurring are considered in [13]. While some of those approach may apply to the blurring, they would not address the issues of the geometric distortions can occur. In addition, that work presumes consistent blurring, which does not occur for vibration blur on rolling shutters.

2. PHOTOHEADS: FAAD AND ATMOSPHERICS/WEATHER

Even after the images are acquired the atmosphere and weather impacts can be critical for FAAD. Studying them is a challenge as it is hard to collect enough data under varying conditions. To address this we designed a specialized experimental setup we call Photo-head. The experimental setup of the initial photo-head is shown below with example images on the right. This "photo-head" data is unique in that it is a well-known set of images (FERET) that were displayed on a special LCD and then re-imaged from approximately 94ft and 182ft. (We are currently implementing another photo-head setup at much greater distances). At these distances we needed a long FOV lens, for which we used Phoenix 500mm zoom lenses (for 35mm cameras), with C-Mount adapters and Panasonic PAL cameras. The marine LCD was 800x600 resolution with 300NITS and a special anti-reflective coating. For display the FERET face images were scaled up for display. As one can see from the examples on the right, which are all from the same subject, the FERET data has a range of inter-pupil distances, poses and contrasts. This re-imaging model allows the system to control pose/lighting and subjects so as to provide the repeatability needed to isolate the effect of long-distance imaging and weather. As one can see, the collection produced images sufficient for identification but with the types of issues, e.g. loss of contrast and variations in size, that one would expect in a realistic long-distance collection. All experiments herein used FaceIt (V4), the commercial face-recognition system from Identix. This algorithm was one of the top performers in the National Face Recognition vendor tests [4].

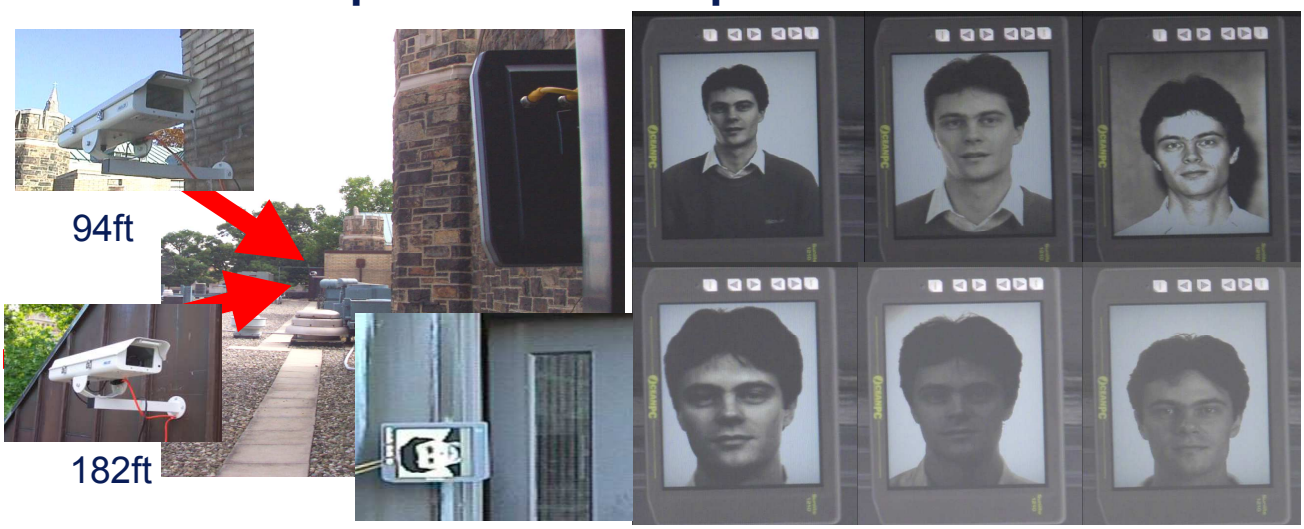

Figure 4: Photohead experimental setup on left, and sample photohead images on the right.

This photo-head dataset is well suited to some of the issues to be encountered in using biometrics for "uncooperative" subjects in surveillance video. One of the most controlled variations is what we call "self-matching" the probe and the gallery are based on the same image, except that the probe has been subject to the distant re-imaging process and the weather. For initial testing we used a camera at approximately 15 ft and the rank-one self-matching performance was over 99%, showing the re-imaging process and LCD are not a significant issue. We then moved to the medium range photo-head collections. We generally ran each probe set, which is 1024 images, with 4 images of each subject, every 15 minuets, with collections over 4 months. The resulting 1.5TB of photo-head data was included in the DARPA HBASE, and subsets of the data are available from the authors. With 4 images per subject we can use the BRR technique [5] to estimate standard errors and statistical confidence. All our graphs include such error bars.

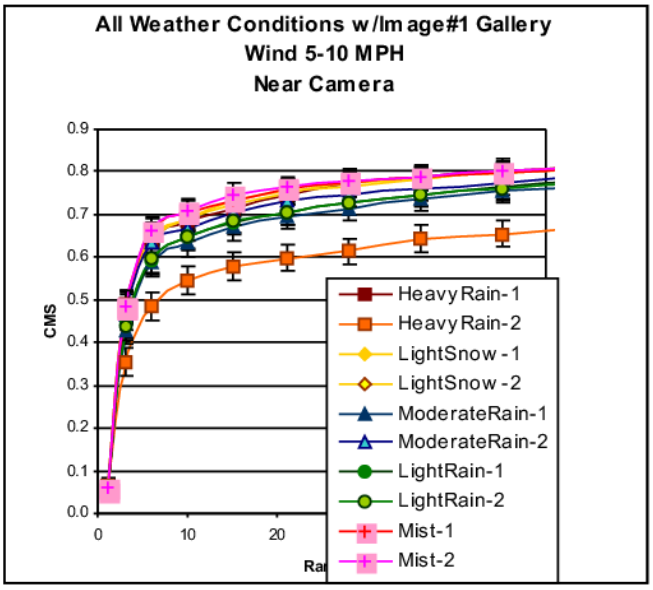

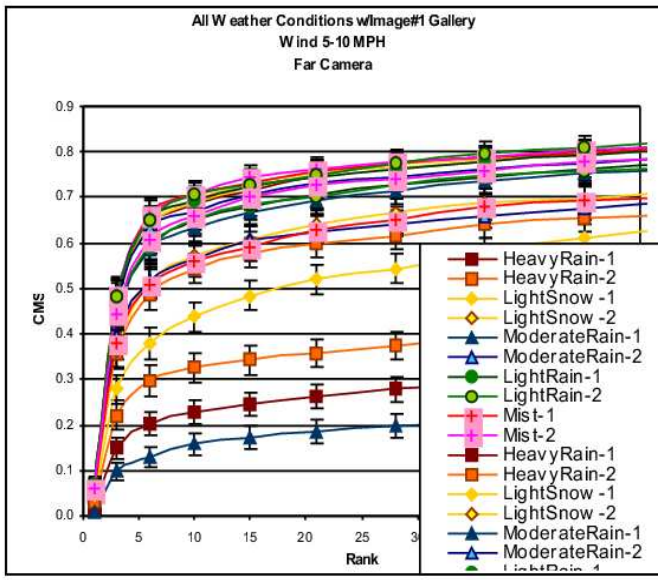

Figure 5: CMC curves for various weather settings.

The two graphs in figure 5 show the impact of different weather conditions on face recognition. These, like most of the results we will show, are cumulative match curves with error bars. The curve shows the recognition rate on the vertical axis and the "rank" used to decide correct recognition on the horizontal axis. Rank-N recognition means the person was within the top N scores of the systems. Two things should be apparent from these graphs. First, looking at rank-1 recognition (or even rank 3), off the shelf systems are simply not sufficient for these ranges, even under the best of weather conditions and ideal pose/expression. (These are self-matching experiments). Second, the far camera, at approximately 182ft, was much more significantly impacted by the variations in weather. (The best weather rank-1 recognition at 182ft was < 60%). While it is not show here, but increasing wind even more significantly impacted the system, in part because at these ranges even a small deflection of the camera causes significant blur and may take the face out of the sensors field of view. (With these long FOV lenses, we needed 30" housings that increase wind loading.) These graphs are computed over more than 20,000 images and with the "controls" of the photo-head collections we know the images are identical, thus the variations are not artifacts of individual errors, pose or expression changes. The techniques of [6] improved performance slightly, not statistically significantly, in large part because they don't address blur or geometric distortion, only contract and dynamic range.

In addition to variations due to obvious weather effects, our experiments also showed that there were variations due to time of day. The error bar for the graphs is shown on the leftmost point of figure 6 is consistent in size for the full graph. These differences are statistically significant. Note that to reduce the impact of pose and lighting variations, these images are using the *exact same image on the display as in the recognition database* – the only variations between the probe and the gallery are those caused by the imaging system. Recall that indoors at 15ft, the performance on this type of data is nearly perfect. Even with this very strong constraint we see that at 182ft on a clear and low-wind day, for Rank N recognition the performance of one of the best commercial algorithms is below 65% with N <4 and still well

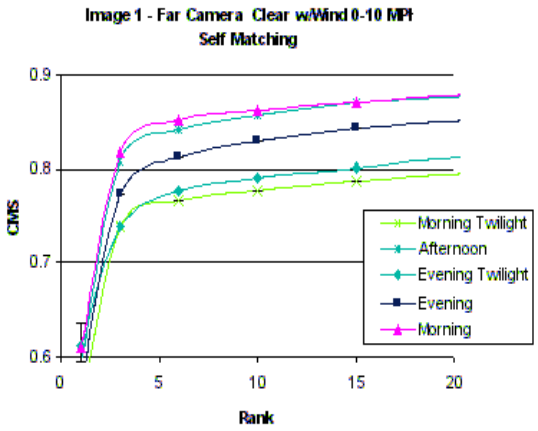

Figure 6: Variations over time of day

Figure 7: Recognition Rank for various "quality" images

below 90% recognition rate even when N is 10. Again, these are averaged over hundreds of trials with 1024 images per trial, so this is not a sampling artifact.

A first guess might be that the impacts were on the raw "image quality'. We examined various measures of facial image quality and (to our surprise) many of the errors had nothing to do with human perceived or algorithmically measured image quality. Figure 7 shows some examples of the recognition rank (i.e. where the image ended up when probes are sorted by match score) for a collection of images from a "same image" experiment. Rank and image quality in this set were inversely correlated.

Our research set off to find the causes of this unexpectedly poor performance. After considerable investigation we hypothesized the poor performance was due in large part to error in localization of the eyes. In [7] we presented an analysis of this theory. To definitively show the cause we added registration markers within our photo-head data to allow us to transform the original eye coordinates to provide eye-locations in the captured images. The graphs above show the recognition performance (with error bars) for the off-the-shelf FaceIt algorithms and when use forced FaceIt to use the correct eye positions. The results on both cameras where statistically significant, and when the eyes are corrected the performance both far and near cameras are similar. These results are, of course, highly optimistic because the data for correction is artificial calibration points and secondly this is self-matching, with the same image as probes and gallery so the near perfect recognition is to be expected.

It is important to note that the "eye-locations" being discussed are not just a question of where in the image the eyes appear but how that position related to where it should be in the image. In the "good quality" images of figure 7, the corrected eye position is not in the middle of the eye! Atmospheric turbulence and lensing effect can distort the face image to the point that to work properly the system needs to use a different eye position for its coordinate system and normalization procedures. Many of the computed eye locations were visibly off the eye, and the average difference between the computed and FaceIt eyes was 6 pixels.

In conclusion, weather and atmosphere have effects that significantly impact the system, as demonstrated in tightly controlled experiments. The impacts grow with distance and is not just a question of loss of contrast or blurring, but also of geometric distortions.

3. FASST: FAILURE FROM SIMILARITY SURFACE THEORY

For non-cooperative subjects at a distance we need some way to determine when the images are good, and when they are poor. At the system level, the simplest application of prediction is in an interactive or on-line system where, if we can predict failure, then we might simply re-acquire a new image and try again. Rather than a pure binary go/no-go, the prediction can be use to determine weights for a temporal fusion module. This has an obvious direct application in biometrics for uncooperative subjects in video. Most commercial systems have some very simple form of this type of failure prediction, e.g. they don't even bother with recognition unless the face detection module reports a face with sufficient confidence. We have introduced a more powerful approach, using the similarity scores that exist after a

potential match. The basic idea is that the shape of the similarity scores viewed as a curve or surface, provides added information about the uniqueness of the match and the probability of success of the overall matching process.

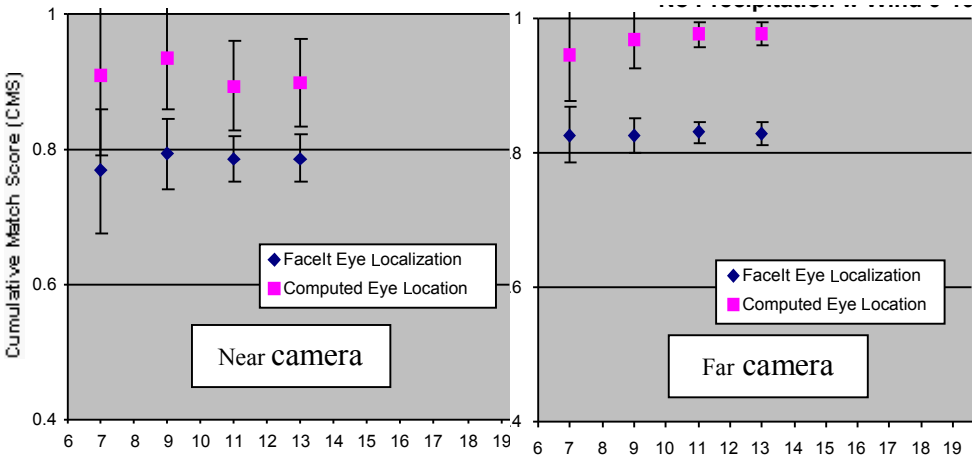

Figure 8: Accuracy with auto detected eyes and with computed "corrected" eyes.

While the general theory suggests that shape analysis should predict failure, the details of the shapes and their potential for prediction are also a function of the data space. Because of the nature of biometric spaces, the similarity surface often contains features at multiple scales caused by matching with sub-clusters of related data, e.g. multiple samples from the same individual over time, from family members, or from people in similar demographic populations. What might be "peaked" in a low-noise system, where the inter-subject variations are small compared to intra-subject variations might be flat in a system with significant inter-subject variations and a large population. These variations are functions of the underlying population, the biometric algorithms, and the collection system. Thus in FASST, the system "learns" the appropriate shape information for a particular system installation. We believe an on-line learning algorithm may be the most appropriate form for final deployment, however initial work will use off-line learning and more careful experimentation.

While multi-sensors might be used in some applications, the goals of the FIINDER™ system will generally be stand-alone or single sensor systems, thus we will be using FASST for fusion from a single sensor, combining different algorithms and perturbations of the original data. We have been exploring two different algorithms to implement FASST-- one based on Ada-Boosting [8] and one based on multi-layer neural nets using wavelet features [9]. Ada boosting builds a strong classifier from multiple rounds of simple combinations of weak classifiers. In essence, it "learns", through a non-linear minimization, the structural combinations of weak classifiers that best account for the training data. It allows for the efficient determination of which features matter for a particular dataset.

This FASST using Ada-boost has been applied to predicting failures of the FaceIT face recognition algorithm by Identix, and tested on both synthetic and natural variations in images. The example on the right shows failure prediction FA/MD rates for both the training and test sets for images collected under real weather conditions. These images were taken at 100ft and 200ft, which make facial recognition far more challenging. The training/testing for this used 21,535 images (split equally). With a FA rate below 2%, the system was still correctly predicting more than 75% of the recognition system failures for one of the leading commercial face-recognition systems. For fusion on fixed data, the FA/MD rate would be quite effective in weighting one algorithm versus another. If a false alarm means taking/processing another image from the video sequence, as it would with most remote recognition of un-cooperative subjects, then an even higher "false alarm" rate can be tolerated. For a miss-detection rate of 10% (i.e. the system correctly predicted 90% of all recognition failures), we have a false alarm rate of only 15%. In summary, the Ada-boosting approach to finding a failure prediction classifier learns the features of importance very quickly, and the resulting classifiers are suitable for real-time operation.

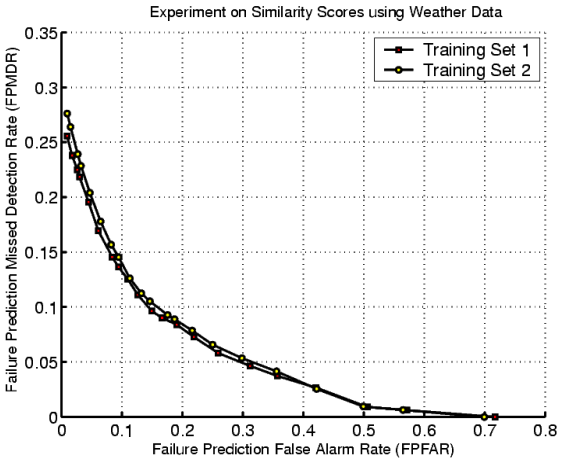

Figure 9: Failure prediction performance for FASST using Adaboost. Tested on 20,000 Photohead images.

The second approach we have been developing uses a back-propagation neural net with wavelet features from the similarity surface. We briefly summarize the approach; details can be found in [9]. This particular approach was developed to address a particular way of improving a biometric system, using failure prediction to predict when a system was likely to fail, and if so, perturbing key input properties to see if better results could be obtained. Note that this example is a special case of biometric fusion, where the fusion is within the same modality.

From our analysis on long-distance images we know that the eye finders in even the leading commercial systems often were inadequate, in part because of atmospheric distortions [9]. So as the first application of the failure prediction theory, we addressed the problem of deciding when a face image was likely to fail, and if so, we considered multiple perturbations of the assumed eye locations and tested the new locations for potential failure.

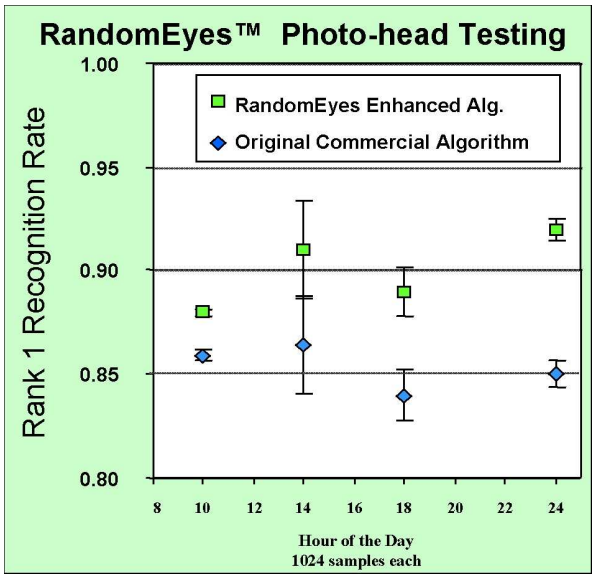

Figure 10: RandomEyes Performance Improvements on a leading commercial algorithm.

The algorithm was tested on the photo-head datasets) using two different algorithms, one of the leading the commercial products, and EGBM as implemented by [3]. The figure on the right shows the overall improvement combining the commercial system with the FASST perturbations. Four different times of day throughout the month of May were used for this analysis. The results are shown averaged over multiple days, grouped by time of day. The error bars show 95% statistical confidence. Note that in this set of experiments, errors in eye localization come from two sources: the eye localization error due to degradation of the input image as a result of atmospheric effects, and the eye localization error due to possible weaknesses in the commercial eye localization algorithm. Together, eye localization error is clearly an unknown quantity, however it is exploited quite effectively here to improve overall classification. It is clear that the improvements were statically significant. Improvements for EGBM were similar, in the 4%-8% range.

4. BIOTOPES® FACES

Recently work by Dr. Boult developed the concept of Biotopes® revocable biometric-based tokens that protect the privacy of the original user, provide for many simultaneous variations that cannot be linked. More significantly, however, we are using Biotopes® faces for this project because the transform induces a robust distance/similarity metric for use in matching in encoded space. This robust distance transform has been shown to significantly improve performance and is expected to be even more significant when used in difficult settings [10]

While many face algorithms have been developed, after finding features most have used a L2 distance measure or a slightly more robust form such as Mahalanobis Cosine to measure distance between probe and gallery. These types of measures have the disadvantage that a few features being significantly off can have a dramatic impact on the recognition rate. The transform and matching used for biotopes basically reduce the penalty for being too far from the particular feature to a constant. Thus features that are significantly in error have a constant (and hence greatly reduced) impact. This can be very important in less than ideal setting for which the FIINDER™ system is being designed because with poor lighting and non-cooperative subjects individual features are more likely to have anomalous measurements.

5. PUTTING IT ALL TOGETHER: FIINDER FOR FAAD

The above three technologies are being combined, with an EMCCD and FPGA computing core, to form the core of the FIINDER system for Face At A Distance. A block diagram is shown in Figure 10. The FPGA-enhanced system does the face-finding in hardware. Because of the lighting issues, which will likely be magnified by the use of the EMCCD, we start with a special lighting normalization described in section 1. The core uses the perturbation approach, a multi-pose gallery, plus multiple algorithms (two variations of the Biotope approach) with a FASST-based prediction fusing the perturbations to account for long-range atmospherics and pose issues of non-cooperative subjects. FASST is also being used for temporal fusion as we did in [11]. We have already developed a face-detection and localization routine that is suitable for FPGA implementation. Mapping this to embedded systems hardware, including the formal fixed-

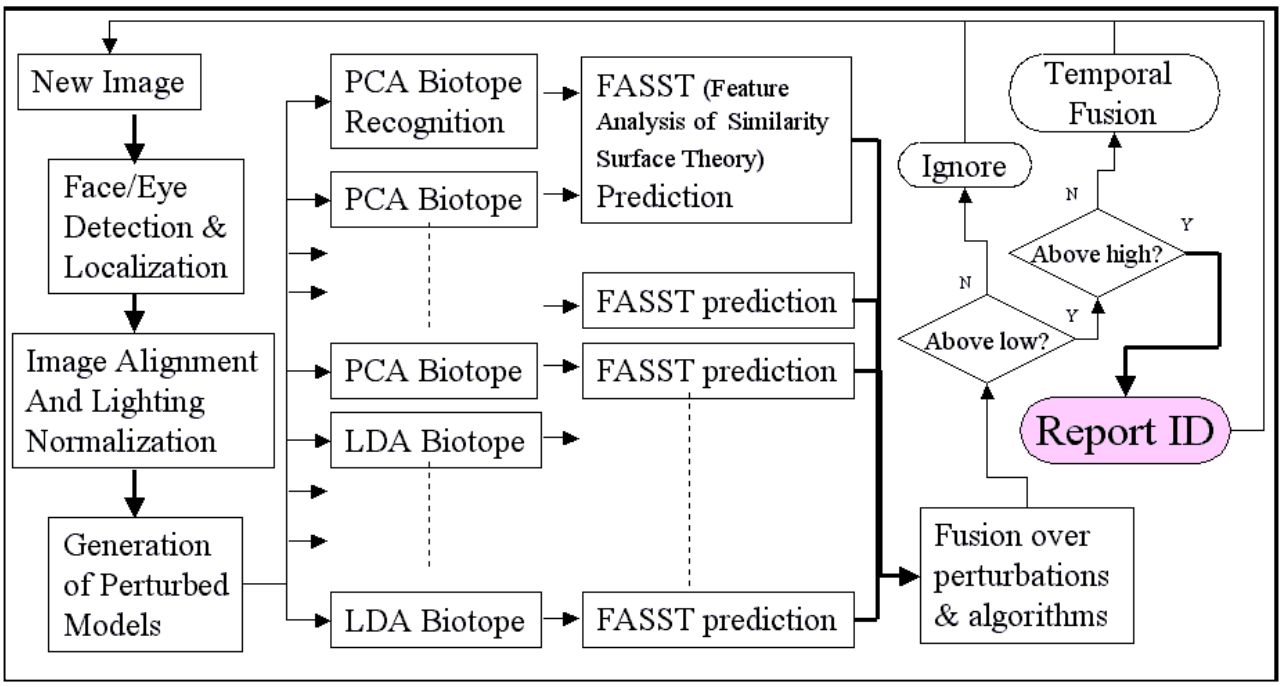

Figure 10: Overall FIINDER architecture for Face At A Distance.

point analysis is underway and that full fixed-point algorithm, with which we will be doing the formal testing. The remaining components exist, and are being ported to the FPGA environment, which will undergo engineering and prototype testing in the field to take the system to TRL6.

An important question is, of course, how well this performs. Unfortunately, no existing dataset is very well suited for proper performance testing of this novel approach and a collection is being designed for Spring 2008. To provide an overall idea of its performance, the techniques have been tested on multiple well known datasets. Table 2 shows results of the revocable Biotopes on 4 FERET datasets. To demonstrate the generality of the improvements to be obtained by using Biotopes and FIINDER we extended algorithms included in the Colorado State University (CSU) Face Identification Evaluation System (Version 5.0)[3]. In particular we developed Biotope versions from the "baseline" PCA-based face recognition system using multiple metrics, from their LDA-based face recognition algorithms and from the Elastic Bunch Graph Matcher (EBGM). As is evident, applying the revocable transforms to convert PCA into a biotope algorithm very significantly improved its performance. (Also for LDA and EBGM). While the best of the original FERET is shown in the table, commercial algorithms have continued to improve since that date. For comparison, the highest reported score for any commercial algorithm on Dup1 is 89%, so the group LDA was as effective as any known commercial algorithm and the "individual" biotope algorithms were better. The "Individual" Biotopes performed better, but require

Algorithm \ Dataset	Dup1	Dup2	Fafb	Fafc
# Matched scores	479	159	1195	194
# Non-matched	228 K	25 K	1427K	37 K
PCA L2	33.7	14.1	74.3	4.6
LDA ldaSoft	44.1	18.8	70.9	41.7
EBGM Predictive	43.6	24.7	86.9	35.5
FERET "BEST"	59.1	52.1	86.2	82.1
Individual PCA Biotope® Face	90.7	87.2	99.5	100
GroupRobust PCA Biotope® Face	86.6	85.5	98.3	100
LDA Biotopes® Face	90.7	87.2	99.5	100
GroupRobust LDA Biotope® Face	88.9	85.5	98.9	100
Individual EBGM Biotope® Face	91.3	88.0	100.	100
Multi-image FIINDER™ PCA	**100**	**100**	**100**	**100**
Multi-image FIINDER™ LDA	**100**	**100**	**100**	**100**
PCA L2 w Securics Normalize	57.2	13.2	59.0	60.2

Table 2: Rank 1 Recognition Rates on FERET subsets

multiple images of an individual for enrollment. The performance of EBGM was superior to the PCA and LDA based algorithms, but the cost and structure of that algorithm make it unacceptable for our real-time system. The final row of table 2 shows the performance gain for the PCA algorithm just from using the new Securics Normalization introduced in Section 1. With improved intensity normalization performance was significantly increased on Dup1 and FAFC, with similar performance on Dup2, but a slight loss of performance on FAFB. Note the "gallery" was the PCA templates of the stock gallery with the standard normalization, and that may have caused the loss of performance on FAFB and Dup2. The table also shows that the multi-image FIINDER approach, implemented in floating point for the individual Biotopes versions, with the new normalization and perturbations, using either PCA or LDA alone achieve 100% recognition accuracy on this data. This 100% recognition rate simply shows the datasets are too small for effective testing of this algorithm Testing with larger datasets is expected after the fixed-point version of the FIINDER system is complete.

The second dataset we used for testing is the Pose Illumination and Expression (PIE) dataset from CMU, [Gross-et-al-04]. Here we considered only the lighting and illumination, not the pose or expressions. The performance for the Multi-image FIINDER™ algorithms was again 100% recognition, but PIE only has 69 subjects so while it shows robustness to lighting, the gallery size is still quite small. In comparison, even with the improved Dual-LUT normalization the standard PCA and LDA algorithms only achieve 83% and 72% rank-1 recognition respectively.

6. CONCLUSIONS

This paper reviewed our work on major issues for Face at a Distance (FAAD), including sensor issues and weather/atmospherics. It reviewed the Photo-head concept and results for systematic evaluation of atmospheric and weather effects. This paper showed techniques that, when combined, addressed these issues, and introduced the FIINDER architecture for effective FAAD. The conclusion is: non-cooperative face-at-a-distance, especially in low light, is a challenging problem that must be addressed at a every level of the system design, including the sensor, the lenses, and the algorithm, with the system predicting its own performance.

7. REFERENCES

[1] D. Socolinsky, L. Wolff and A. Lundberg, "Image Intensification for Low-Light Face Recognition", IEEE Workshop on Biometrics June 2006, in combination with CVPR2006.

[2] T. L. Vogelsong, T. E. Boult, D. W. Gardner, R. Woodworth, R. C. Johnson, and B. Heflin, "24/7 Security System: 60 FPS Color EMCCD Camera with Integral Human Recognition", SPIE Defense Symposium: Sensors, and Command, Control, Communications, and Intelligence (C3I) Technologies for Homeland Security and Homeland Defense VII, April 2007.

[3] Bolme, D.S., Beveridge, J.R., Teixeira, M. and Draper, B.A. (2003). "The CSU face identification evaluation system: its purpose, features, and structure". ICVS 2003: 304-313

[4] Phillips, P.J., Grother, P., Micheals, R.J., Blackburn, D.M., Tabassi, E., and Bone, M. (2003). Face Recognition Vendor Test 2002 (FRVT 2002). Technical Report NISTIR 6965, NIST.

[5] R.J. Micheals, T.E. Boult, "Efficient evaluation of classification and recognition systems", IEEE Conf. on Computer Vision and Pattern Recognition (CVPR) pp. 50-57, June 2001.

[6] S. G. Narasimhan and S. K. Nayar , "Contrast Restoration of Weather Degraded Images", IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI), Vol. 25, No. 6, June 2003

[7] T.P. Riopka and T.E. Boult, The Eyes Have It. ACM Biometrics Methods and Applications Workshop, Berkeley, CA. pp. 33-40. 2003.

[8] W. Li, X. Gao and T. E.Boult, "Predicting Biometric System Failure", IEEE Conference on Computational Intelligence for Homeland Security and Personal Saftey, March 2005.

[9] T.P. Riopka and T.E. Boult, Classification Enhancement via Biometric Pattern Perturbation, IAPR Conference on Audio- and Video-based Person Authentication (AVBPA), July 2005.

[10] T.E. Boult, ""Robust distance measures for face recognition supporting revocable biometric tokens", IEEE Conf. on Face and Gesture, April 2006

[11] B. Xie, V. Ramesh, Y. Zhu Terry Boult, "On Channel Reliability Measure Training for Multi-Camera Face Recognition", IEEE Workshop on the Application of Computer Vision, Feb 2007

[12] D. A. Socolinsky, .Selinger: "Thermal Face Recognition in an Operational Scenario", IEEE CVPR 2004:1012-1019

[13] Y. Yao, B. Abidi, N. D. Kalka, N. Schmid, and M. Abidi, High Magnification and Long Distance Face Recognition: Database Acquisition, Evaluation, and Enhancement, in Proc. Biometrics Consortium Conf., 2006