

Thermal Face Recognition in an Operational Scenario *

Diego A. Socolinsky† Andrea Selinger‡

†Equinox Corporation ‡Equinox Corporation
207 East Redwood St 9 West 57th St
Baltimore, MD 21202 New York, NY 10019

{diego, andrea}@equinoxsensors.com

Abstract

We present results on the latest advances in thermal infrared face recognition, and its use in combination with visible imagery. Previous research has shown high performance under very controlled conditions, or questionable performance under a wider range of conditions. This paper shows results on the use of thermal infrared and visible imagery for face recognition in operational scenarios. In particular, we show performance statistics for outdoor face recognition and recognition across multiple sessions. Our results support the conclusion that face recognition performance with thermal infrared imagery is stable over multiple sessions, and that fusion of modalities increases performance. As measured by the number of images and number of subjects, this is the largest ever reported study on thermal face recognition.

1 Introduction

Over the last few years, there has been a surge of interest in face recognition using thermal infrared imagery. While the volume of literature on the subject is notably smaller than related to visible face recognition, there is nonetheless a steady stream of research [1, 2, 3, 4, 5, 6]. Previous work centered mostly on validating infrared imaging as a viable tool for biometric identification. These studies relied on databases limited both in size and variability, due to the expense and complexity of extensive data collection. Early results were based on gallery and probe sets collected indoors during a single session. In that respect, they resemble the **fa/fb** tests in the FERET program [7].

Comparable performance for visible and thermal face recognition was reported in [3], using a small database of

imagery collected in a single session. Their thermal sensor was a low-sensitivity, low-resolution ferro-electric sensor. In [8, 1], superior performance of thermal imagery is reported using a variety of algorithms. These studies used a database of coregistered visible/thermal image pairs of approximately 90 subjects, collected indoors during a single session. During data collection illumination conditions were purposely varied in order to present a challenge for visible face recognition. Results of a recent time-lapse recognition experiment were reported in [4, 9]. This study uses a database of 240 subjects collected over a 10 week period. Recognition performance was evaluated using a PCA algorithm for both visible and thermal images. The most interesting conclusion of the study is that face recognition using thermal images degrades more sharply than with visible images when probe and gallery are chosen from different sessions. This has obvious negative implications for the use of thermal imagery in face recognition, as any imaginable application of face recognition would require enrollment and testing images to be acquired at different times, and potentially different locations. An additional conclusion of the study in [4, 9] is that despite the degraded thermal recognition performance, fusion of both visible and thermal modalities yields better overall performance.

The current paper sets out to expand the knowledge on visible/thermal face recognition by extending the operational scenario to outdoor imaging. This is well known to be a challenging condition for all existing face recognition systems, and has been highlighted as a critical area of research [10]. We will present results under realistic testing conditions, with gallery and probe images acquired during different sessions, as described in Section 2.

Additionally, we will expand the treatment of time-lapse performance using thermal imagery given in [4, 9]. We will show that while the results in [4, 9] are valid (and indeed reproducible on our data), they are not necessarily a reflection of the modality, but rather of the algorithm used to measure the quality of that modality.

*This research was supported in part by the DARPA Human Identification at a Distance (HID) program, contract # DARPA/AFOSR F49620-01-C-0008.

Report Documentation Page

Form Approved
OMB No. 0704-0188

Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

1. REPORT DATE 2004		2. REPORT TYPE		3. DATES COVERED 00-00-2004 to 00-00-2004	
4. TITLE AND SUBTITLE Thermal Face Recognition in an Operational Scenario				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Equinox Corporation, 9 West 57th Street, New York, NY, 10019				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES The original document contains color images.					
14. ABSTRACT see report					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES 8	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

Our study is the largest ever reported for thermal face recognition, in terms of number of images and number of subjects. In addition, this is the first ever study to consider outdoor and indoor imaging conditions for thermal imaging, and one of few to do so even for visible face recognition. Due to the operational realism of this study, our government sponsor has requested that limited information be disseminated into the public domain as to specific details of the experimental setup and its location. Therefore, most discussion about experimental logistics will be restricted to number of subjects and imaging conditions.

2 Data Collection and Preprocessing

The majority of the imagery used in this study was collected during during eight separate day-long sessions spanning a two week period. A total of 385 subjects participated in the collection. Four of the sessions were held indoors in a room with no windows and carefully controlled illumination. Subjects were imaged against a plain background some seven feet from the cameras, and illuminated by a combination of overhead fluorescent lighting and two photographic lights with umbrella-type diffusers positioned symmetrically on both sides of the cameras and about six feet up from the floor. Due to the intensity of the photographic lights, the contribution of the fluorescent overhead lighting was small. Three of the four indoor sessions were held in different rooms. The remaining four sessions were held outdoors at two different locations. During the four outdoor sessions, the weather included sun, partial clouds and moderate rain. All illumination was natural; no lights or reflectors were added. Subjects were always shaded by the side of a building, but were imaged against an unconstrained natural background which included moving vehicles, trees and pedestrians. Even during periods of rain, subjects were imaged outside and uncovered, in an earnest attempt to simulate true operational conditions. For each individual, the earliest available video sequence in each modality is used for gallery images and all subsequent sequences in future sessions are used for probe images.



Figure 1: Example visible images of a subject from indoor and outdoor sessions.

For all sessions, subjects were cooperative, standing about seven feet from the cameras, and looking directly at them when so requested. On half of the sessions (both indoors and outdoors), subjects were asked to speak while being imaged, in order to introduce some variation in facial expression into the data. For each subject and session, a four second video clip was collected at ten frames per second in two simultaneous imaging modalities. We used a sensor capable of acquiring coregistered visible and longwave thermal infrared (LWIR) video. The visible component has a spatial resolution of 640×480 pixels, and 8 bits of spectral resolution. The thermal sensor is uncooled, and has 12 bits of depth, sensing between 8μ and 12μ at a resolution of 320×240 pixels.



Figure 2: Automatic detection of the face and eyes shown on an overlay of visible and thermal images. (The lower cross-hairs denotes the center of the face, not the nose)

Faces were automatically detected in all acquired indoor and outdoor frames, using a system based on the algorithm described in [11]. No operator intervention was required for this step, the results of which are shown in Figure 2 on an overlaid visible/thermal representation. The same figure shows the results of eye localization, which was also performed automatically on every frame. Recall that since visible and thermal images are coregistered, eye locations in one modality give us those in the other. The detected eye locations were used to affinely transform all images to a standard grid of 99×132 pixels, with fixed eye locations, with all necessary interpolation done bilinearly. Once geometrically normalized, all images were masked to exclude background. We should emphasize the fact that all data used for the experiments below was processed in a completely automatic fashion, once again in an attempt to simulate true operational conditions.

Thermal images in this study were processed via one-point calibration in order to compensate for non-uniformities in the microbolometer array. This simply consists of subtracting from each image pixel the response of that pixel to a uniform source. More details on calibration of thermal imaging sensors can be found in [8].



Figure 3: Original (left) and illumination compensated (right) outdoor images.

Visible images were run through a simple procedure in order to eliminate some of the most severe effects of outdoor illumination. Given the shape of human heads, self shadowing of one side of the face during strong lighting conditions is a major source of appearance variation. As long as there is sufficient dynamic range in the image, this problem can be attenuated through the following process. We estimate the means and variances of the pixels on either side of the face and use them in a simple criterion to determine the better illuminated side. We rescale the pixels on the poorly illuminated side to the mean and standard deviation of the good side. A sharp transition between both sides of the processed face is avoided by combining them through a weighted average near the centerline.

This simple technique is quite effective for compensating for common lateral self shadowing, as can be seen in the results below, but does not help with shadowing of the eye sockets from the superciliary arches, which is very common with strong overhead illumination from the sky or sun. Also, overexposure from excessive illumination is somewhat common outdoors, where the dynamic range of lighting is considerably larger than indoors. In these cases, we use another heuristic procedure which we have found quite effective. We note that the skewness of the distribution of grayvalues of an underexposed image is larger than that of a well illuminated image, which is in turn

larger than that of an overexposed image. Also, note that gamma-correction with an exponent larger than unity decreases skewness while the opposite is true for an exponent below unity. Combining these two observations, we use a gamma-correction step with an exponent dependent on the skewness of the distribution of grayvalues on the face. Figure 3 shows the effect of this process for two outdoor visible images. We see in Section 5, Figure 5, that this preprocessing step has a favorable effect on both PCA- and LDA-based recognition.¹

3 Structure of the Experiments

We performed experiments with three different algorithms in each of the two modalities: PCA with Mahalanobis angle distance, LDA with angle distance and the (blinded for review) algorithm. The first two are standard algorithms with performance evaluations widely available in the literature, including [2], in which the authors present a comprehensive analysis of their performance on visible and thermal infrared imagery in a same-session recognition scenario. The third one is a commercial algorithm made available for testing in binary form.²

The training set for all algorithms was completely disjoint from gallery and probe images, in time, space and subjects. That is, the training set was collected at an earlier time, in a different location and used a disjoint set of subjects. This insures that the results reported below are indicative of real world performance. Since the data collection involved video data in both modalities, we evaluated recognition performance using short video sequences as input. Following the recent trend in evaluation of biometric algorithms [12, 13], we performed randomized experiments to estimate both mean recognition rates and confidence intervals for all tests. Each test (regardless of modality or algorithm) used a random sampling of three images from the gallery sequence of each individual and four consecutive frames from a random probe sequence of each individual, with a random starting frame within the sequence. The distance from a probe sequence to an individual in the gallery was defined to be the smallest distance between any frame in the sequence and any image of that individual in the gallery. Classification was based on nearest neighbors with respect to this distance. For each modality and algorithm we performed one-hundred random repetitions of the experiment, using the same sampling pattern for all algorithms and modalities. We then computed the mean recognition rate at each rank, from one to ten, along with the standard deviation of that measurement over the one-hundred

¹Illumination preprocessing has no measurable effect on the remaining algorithm.

²This algorithm was made available for testing purposes at [http://\(blinded for review\)](http://(blinded for review)).

trials. All performance graphs below depict average performance over the whole trial, with error bars corresponding to 95% confidence intervals (or equivalently, 1.96 standard deviations).

We report results for fusion of visible and thermal imagery for all algorithms and modalities. Following [9], we assign a score to a visible-thermal image pair that is the sum of the scores of each image in the pair. This addition is done with equal weights. When we report results for fusion below, we refer to the performance resulting from using a nearest neighbor classifier on the sum of scores.

4 Thermal Infrared Phenomenology

While the nature of face imagery in the visible domain is well-studied, particularly with respect to illumination dependence [14], its thermal counterpart has received less attention. In [4], the authors show some variability in thermal emission patterns during time-lapse experiments, and properly blame it for decreased recognition performance. Figure 4 shows comparable variability within our data. The left column shows enrollment images and the right column shows test images from the same subject at a later session. We can plainly see how emission patterns are different around the nose, mouth and eyes. Weather conditions during our data collection were quite variable, with some days being substantially colder and windier than others. In addition, some subjects were imaged indoors immediately after coming from outside, while others had as much as twenty minutes of waiting time indoors before being imaged. These conditions contribute to a fair amount of variability in the thermal appearance of the face. When exposed to cold or wind, capillary vessels at the surface of the skin contract, reducing the effective blood flow and thereby the surface temperature of the face. When a subject transitions from a cold outdoor environment to a warm indoor one, a reverse process occurs, whereby capillaries dilate, suddenly flushing the skin with warm blood in the body’s effort to regain normal temperature.

Additional fluctuations in thermal appearance are unrelated to ambient conditions, but are rather related to the subject’s metabolism. During our data collection, an uncontrolled portion of the subjects engaged in strong physical activity at different periods prior to imaging. The time elapsed from physical exertion to imaging was uncontrolled and known to be different for different sessions. This further contributes to the change in thermal appearance. Also, high temporal frequency thermal variation is associated with breathing. The nose or mouth will appear cooler as the subject is inhaling and warmer as he or she exhales, since exhaled air is at core body temperature, which is several degrees warmer than skin temperature.



Figure 4: Variation in facial thermal emission from two subjects in different sessions. Left column is the enrollment image and right column is the test image.

Much like recognition from visible imagery is affected by illumination, recognition with thermal imagery is affected by a number of exogenous and endogenous factors. And while the appearance of some features may change, their underlying shape remains the same and continues to hold useful information for recognition. Thus, much like in the case of visible imagery, different algorithms are more or less sensitive to image variations. As we see in Figure 5, for example, proper compensation for illumination prior to recognition has a favorable effect on recognition performance with visible imagery. Clearly, the better algorithms for thermal face recognition will perform equivalent compensation on the infrared imagery prior to comparing probe and gallery samples.

5 Experimental Results and Discussion

We performed all experiments as described in Section 3. Enrollment images for all experiments were taken from indoor sessions since this is the most likely scenario for an access control system: users are enrolled in an office at the same time that they are issued their identification cards, and they later seek access at a different location, either indoors




Figure 5: Cumulative recognition rates for all algorithms and conditions. Left: visible imagery without illumination compensation. Center: visible imagery with illumination compensation. Right: LWIR imagery.




Figure 6: Recognition results by algorithm for indoor enrollment and indoor testing. Note that the vertical scales are different in each graph. Left: PCA with illumination compensation. Center: LDA with illumination compensation. Right: (blinded for review) algorithm




Figure 7: Recognition results by algorithm for indoor enrollment and outdoor testing. Note that the vertical scales are different in each graph. Left: PCA with illumination compensation. Center: LDA with illumination compensation. Right: (blinded for review) algorithm

or outdoors. Two sets of experiments are presented, those with test imagery acquired indoors and outdoors. The ones with outdoor test imagery are easily representative of an access control point situated at the entrance of a building or at a roadside checkpoint. Indoor test images were acquired with very structured illumination, and are therefore easier (at least for the visible half) than should be expected for an indoor access control point.

A summary of top-match recognition performance is shown in Tables 1 and 2. A quick glance yields some preliminary observations. Under controlled indoor conditions, two of the visible algorithms are probably showing saturated performance on the data, which indicates that the test is too easy according to [15]. This may also be the case for the best thermal algorithm. Across the board, for both indoor and outdoor conditions, fusion of both modalities improves performance over either one separately. Comparing indoor versus outdoor performance shows that the latter is considerably lower with visible imagery, and significantly so even with thermal imagery. Fusion of both modalities improves the situation, but performance outdoors is statistically significantly lower than indoors, even for fusion. This difference, however, is much more pronounced for the lower performing algorithms, which is simply a reflection of the fact that the better algorithms have superior performance with more difficult data, without sacrificing performance on the easy cases.

Figure 5(left) shows the marked improvement that illumination compensation, as described in Section 2 has on visible recognition performance. We additionally experimented with the symmetric shape-from-shading method in [16] and found that our simple preprocessing yielded better results. Since the improvement is so large, all results below for visible imagery include illumination compensation. In Figure 5(center) we see a side-by-side comparison of visible recognition performance for all algorithms under indoor and outdoor conditions. In this case, the ordering of the algorithms in terms of performance is the same indoors and outdoors, with all outdoor results underperforming all indoors results. This clearly indicates that even when attempting to compensate for severe outdoor illumination, the variability induced by imaging conditions overpowers intrapersonal similarity. For indoor imagery with carefully controlled illumination, the top two algorithms are extremely close, and both of them have very good performance. For outdoor conditions, all algorithms are statistically different, but even the best performer only reaches 67% top-match recognition.

Results for all algorithms using thermal infrared imagery are shown in Figure 5(right). It is interesting to note that in this case the results are ordered by algorithm, rather than imaging conditions, and all differences in top-match recognition performance are statistically significant. This pre-

	Vis	LWIR	Fusion
PCA	81.54	58.89	87.87
LDA	94.98	73.92	97.36
(blinded for review)	97.05	93.93	98.40

Table 1: Top-match recognition results for indoor probes

	Vis	LWIR	Fusion
PCA	22.18	44.29	52.56
LDA	54.91	65.30	82.53
(blinded for review)	67.06	83.02	89.02

Table 2: Top-match recognition results for outdoor probes

sumably indicates that the variation induced by the imaging conditions is smaller than in the visible case. Performance results for indoor recognition experiments by algorithm are shown in Figure 6. Note how performance with visible imagery is relatively close among algorithms, which is not surprising since the imagery is very carefully controlled. More interestingly, we see that while the difference in performance between visible and thermal imagery is very significant for PCA and LDA at top-match and remains so for the top ten, it is barely statistically significant at top-match when using the (blinded for review) algorithm, and that significance vanishes by the third rank. This indicates that this algorithm compensates for some sources of intrapersonal variability which the other two do not. Also, this raises the issue of how to evaluate the usefulness of an imaging modality for a specific task, in this case face recognition. If we look at the results using PCA, as in [4, 9], we would rightly conclude that there is a severe loss of performance associated with the use of thermal imagery with respect to visible imagery, at least when the illumination is carefully controlled. However, we see that if we measure performance using another algorithm, that loss of performance may be much smaller, or even vanish. Therefore, we must keep in mind that when we judge the value of an imaging modality for a given task, we must try to separate algorithmic effects from intrinsic value. This is not easily done, however, since we can only measure the value of the outcome, and not the modality itself.

Looking at the results from the outdoor experiments in Figure 7, we see clear indication of the difficulty of outdoor face recognition with visible imagery. All algorithms have a difficult time in this test, and even the best performer achieves only about 84% recognition at rank 10. Thermal performance is also lower for all methods than with indoor imagery, but not so much as in the visible case. However, in this case the performance difference between the modalities is very significant for all three algorithms. It is clear from

this experiment, as from those in [10] that face recognition outdoors with visible imagery is far less accurate than when performed under fairly controlled indoor conditions. For outdoor use, thermal imaging provides us with a considerable performance boost.

Fusion of both imaging modalities improves performance under all tests and algorithms, even when using the simple combination rule described in Section 3. This supports previous results reported in [2, 4]. As we mentioned above, it is interesting to note that while even for the best performing algorithm there is a statistically significant difference between fusion performance outdoors and indoors, that significance is smaller the better the algorithm. This is a reflection of the fact that all methods perform well with easy data, but only the better methods perform well in difficult conditions.

6 Conclusion

We presented visible and thermal face recognition results in an operational scenario including both indoor and outdoor settings. Our study is the first ever to consider outdoor and indoor imaging conditions for thermal imaging, and one of few to do so even for visible face recognition. With images of 385 subjects collected over a two-week period, it is also the largest ever reported for thermal face recognition in terms of number of images and number of subjects.

Every effort was made to produce a study that would properly reflect the performance of face recognition technology, both visible and thermal, in a real-world application. To that effect, the training set used for all algorithms was collected at an earlier time, in a different location and used a disjoint set of subjects. Additionally, and unlike most published results, all feature detection and image normalization was done automatically, without manual intervention. This included detection of the face itself and of the facial landmarks necessary for geometric alignment. As a result, our recognition rates are likely to be representative of expected field outcomes. The statistical significance of our analysis is based on our randomized approach to selecting gallery and probe images for experiments with three different algorithms in each of the two modalities: PCA with Mahalanobis angle distance, LDA with angle distance and the (blinded for review) algorithm.

While the visible imagery was affected by changes in illumination outdoors, thermal imagery was affected both indoors and outdoors by a number of factors such as physical exertion of subjects and weather conditions, resulting in better performance obtained with visible imagery indoors under controlled lighting conditions. However, the performance difference between modalities varies depending on the algorithm. This leads us to remark that while evaluating the suitability of an imaging modality for a specific task by

comparing outcomes of a given algorithm is a reasonable surrogate, we must realize that we are measuring the joint value of the algorithm and the modality. It is difficult to decouple the two effects, but at the very least we must be aware of the connection. This was particularly relevant for our study, where we noticed that while multi-session thermal face recognition under controlled indoor illumination was statistically poorer than visible recognition with two standard algorithms, significance was substantially reduced with an algorithm more specifically tuned to thermal imagery. This suggests that previous results reported on multi-session thermal face recognition may be incomplete.

Outdoor recognition performance is worse for both modalities, with a sharper degradation for visible imagery regardless of algorithm. It is clear from our experiments that face recognition outdoors with visible imagery is far less accurate than when performed under fairly controlled indoor conditions. For outdoor use, thermal imaging provides us with a considerable performance boost. Thermal recognition performance suffers a moderate decay when performed outside against an indoor enrollment set, probably as a result of environmental changes. As previously reported, fusion of both imaging modalities improves performance under all tests and algorithms, even when using a simple combination rule. This improvement is particularly relevant outdoors, where performance of each individual modality is impaired. In fact, fused performance outdoors is nearing the levels of indoor face recognition, making it an attractive option for human identification in unconstrained environments.

References

- [1] D. A. Socolinsky and A. Selinger, "A comparative analysis of face recognition performance with visible and thermal infrared imagery," in *Proceedings ICPR*, Quebec, Canada, August 2002.
- [2] D.A. Socolinsky and A. Selinger, "Face recognition with visible and thermal infrared imagery," *Computer Vision and Image Understanding*, July - August 2003.
- [3] Joseph Wilder, P. Jonathon Phillips, Cunhong Jiang, and Stephen Wiener, "Comparison of Visible and Infra-Red Imagery for Face Recognition," in *Proceedings of 2nd International Conference on Automatic Face & Gesture Recognition*, Killington, VT, 1996, pp. 182–187.
- [4] X. Chen, P. Flynn, and K. Bowyer, "Visible-light and infrared face recognition," in *Proceedings of the Workshop on Multimodal User Authentication*, Santa Barbara, CA, December 2003, to appear.
- [5] B. Abidi, "Performance comparison of visual and thermal signatures for face recognition," in *The Biometrics Consortium Conference*, Arlington, VA, September 2003.
- [6] F. J. Prokoski, "History, Current Status, and Future of Infrared Identification," in *Proceedings IEEE Workshop on*

Computer Vision Beyond the Visible Spectrum: Methods and Applications, Hilton Head, 2000.

- [7] P. Jonathon Phillips, Hyeonjoon Moon, Syed A. Rizvi, and Patrick J. Rauss, "The FERET Evaluation Methodology for Face-Recognition Algorithms," Tech. Rep. NISTIR 6264, National Institute of Standards and Technology, 7 Jan. 1999.
- [8] D. Socolinsky, L. Wolff, J. Neuheisel, and C. Eveland, "Illumination Invariant Face Recognition Using Thermal Infrared Imagery," in *Proceedings CVPR*, Kauai, Dec. 2001.
- [9] X. Chen, P. Flynn, and K. Bowyer, "PCA-based face recognition in infrared imagery: Baseline and comparative studies," in *International Workshop on Analysis and Modeling of Faces and Gestures*, Nice, France, October 2003.
- [10] P. J. Phillips, D. Blackburn, M. Bone, P. Grother, R. Micheals, and E. Tabassi, "Face recognition vendor test 2002 (FRVT 2002)," Available at <http://www.frvt.org/FRVT2002/default.htm>, 2002.
- [11] D. Socolinsky, J. Neuheisel, C. Priebe, and J. DeVinney, "A boosted ccd classifier for fast face detection," in *35th Symposium on the Interface*, Salt Lake City, UT, March 2003.
- [12] R. J. Micheals and T. Boulton, "Efficient evaluation of classification and recognition systems," in *Proceedings of CVPR*, Kauai, HI, December 2001.
- [13] J. R. Beveridge, K. She, B. A. Draper, and G. H. Givens, "Parametric and Nonparametric Methods for the Statistical Evaluation of Human ID Algorithms," in *Proceedings of the Third Workshop on Empirical Evaluation Methods in Computer Vision*, Kauai, HI, Dec. 2001.
- [14] Yael Adini, Yael Moses, and Shimon Ullman, "Face Recognition: The Problem of Compensating for Changes in Illumination Direction," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 19, no. 7, pp. 721–732, July 1997.
- [15] J. Phillips and E. Newton, "Meta-analysis of face recognition algorithms," in *Proceedings of IEEE International Conference on Automatic Face and Gesture Recognition*, Washington D.C., May 2002.
- [16] W. Zhao and R. Chellappa, "Robust face recognition using symmetric shape-from-shading," Tech. Rep., Center for Automation Research, University of Maryland, College Park, 1999, Available at <http://citeseer.nj.nec.com/zhao99robust.html>.