

Recent research results in iris biometrics

Karen Hollingsworth, Sarah Baker, Sarah Ring
Kevin W. Bowyer, and Patrick J. Flynn

Computer Science and Engineering Department,
University of Notre Dame, Notre Dame, IN 46556

ABSTRACT

Many security applications require accurate identification of people, and research has shown that iris biometrics can be a powerful identification tool. However, in order for iris biometrics to be used on larger populations, error rates in the iris biometrics algorithms must be as low as possible. Furthermore, these algorithms need to be tested in a number of different environments and configurations. In order to facilitate such testing, we have collected more than 100,000 iris images for use in iris biometrics research. Using this data, we have developed a number of techniques for improving recognition rates. These techniques include fragile bit masking, signal-level fusion of iris images, and detecting local distortions in iris texture. Additionally, we have shown that large degrees of dilation and long lapses of time between image acquisitions negatively impact performance.

Keywords: iris recognition, iris biometrics, fragile bits, signal-level fusion, pupil dilation, time lapse

1. INTRODUCTION

Iris biometrics research is an active and growing field.¹ Our research in iris biometrics can be divided into two categories: (1) efforts to improve recognition accuracy and (2) efforts to evaluate performance across varying image acquisition conditions. Efforts in both categories are vital to the future iris biometrics applications. Our efforts in the first category include masking fragile bits in the iris code (Section 3), fusing multiple iris images together (Section 4), and detection of local texture distortions (Section 5). Our efforts in the second category include measuring performance under different lighting conditions (Section 6) and measuring performance when images are taken years apart (Section 7).^{*} To enable such research, the University of Notre Dame collects vast databases of biometric information including both still iris images and iris videos. Section 2 gives details on the data sets and software that we have created.

2. DATA SETS AND SOFTWARE

Our research group has been acquiring iris image data for use in government evaluation programs and our own research since 2004. Recent data collections have used newer sensors, capture and storage of iris videos, and capture of the entire face (including the irises) under near-infrared illumination. To date, in excess of 100,000 iris still images and several thousand iris videos of various types have been collected.

The primary sensors used in our data collection include the following

- **LG 2200** – This is a first-generation iris camera that employs an analog (NTSC format) CCD sensor, three illuminants, and external (PC-hosted) digitization hardware accompanied by driver software that controls the sensor to optimize image quality. The LG2200's analog output is available and can be digitized to acquire video. The image datasets that we have acquired with the LG 2200 intentionally include a broader range of image qualities than that system would normally retain for routine use.⁷
- **LG 4000** – This is a third-generation iris camera that includes three sensors. Two sensors are NIR-sensitive and they are used to acquire iris images of the left and right eye simultaneously. The third camera is a color imager that can be configured to capture a face image. No analog capture capability is available and digital video capture from this camera is not possible with standard software.

Corresponding author: Kevin W. Bowyer, kwb@cse.nd.edu

^{*}Portions of the work reported in this paper are described in greater detail in our other publications.²⁻⁶

- **IOM** – The IOM (Iris On the Move) system is a prototype developed by Sarnoff Research.⁸ It captures videos of the face (including eyes and irises) as the subject walks through a portal that contains an NIR illuminant array. The IOM has multiple cameras to accommodate human subjects of different height ranges, and each camera produces a 2048×2048 video sequence with between five and twenty frames.

The research data sets created by our research group that contain iris data include the following.

- **ICE 2005** – The ICE 2005 data set⁹ was captured in Spring 2004, released in Fall 2005 and consists of 2953 iris images captured from 132 human subjects.
- **ICE 2006** – The ICE 2006 data set⁷ was captured between Spring 2004 and Spring 2005. It is a sequestered data set captured to support a vendor evaluation of iris recognition technology. It contains 59,558 images from 240 human subjects.
- **MBGC Portal Challenge Data 1.0** – this data set was released as part of the first phase of the Multiple Biometrics Grand Challenge evaluation¹⁰ in 2008. It consists of 326 iris video sequences from the LG2200 camera, 582 sequences from the IOM, and 2568 images from the LG2200 or LG4000.

Iris biometrics software based on Masek’s Matlab implementation was implemented in C++ and the source code made available to participants in the ICE program. A later version of this software, with additional refinements for improved segmentation¹¹ and masking of fragile bits² (see section 3), is used in much of our iris biometrics research.

3. CONSISTENT AND FRAGILE IRIS CODE BITS

Many iris recognition algorithms are based on the work of John Daugman.¹² Once the iris is located in the image, the annular iris region is “unwrapped”: the raw image coordinates, (x, y) are converted into normalized polar coordinates, (r, θ) , where r ranges from 0 to 1, and θ ranges from 0 to 2π . A texture filter is applied to the normalized image. The complex filter responses are quantized to create a binary iris code. Each complex number is represented as two bits in the iris code. The first bit is a 1 if the real part of the number is positive, and 0 otherwise; similarly, the second bit is a 1 if the imaginary part of the number is positive, and 0 otherwise. In the subsequent matching step, two such binary iris codes are compared, and a decision is based on the fractional Hamming distance (the fraction of bits that disagree).

Algorithms that follow this pattern produce templates in which not all bits have equal value.² Specifically, complex filter responses near the axes of the complex plane produce fragile bits in the iris code: a small amount of noise in the iris image can shift that filter response from one quadrant to the adjacent quadrant, causing the corresponding bit in the iris code to flip. This type of bit is defined as “fragile”; that is, there is a substantial probability of it ending up a 0 for some images of the iris and a 1 for other images of the same iris.

Figure 1 shows an example of the distribution of 54 complex numbers from 54 different images of the same iris. Each of these complex numbers is associated with the same location in the iris code. In particular, this location on the iris code had a highly inconsistent imaginary bit and a highly consistent real bit. As expected, the complex numbers associated with these two bits of the iris code lay close to the positive real axis.

To improve iris recognition accuracy, we can identify and mask fragile bits using the following strategy. If the complex coefficient has a real part very close to 0, mask the corresponding real bit in the iris code. If the complex coefficient has an imaginary part very close to 0, mask the corresponding imaginary bit. In implementing this strategy, we chose to mask bits corresponding to the 25% of complex numbers closest to the axes.

We tested this strategy on a data set of 1226 iris images from 24 subjects.² These images were hand-selected to be mostly free of occluding eyelids and lashes. The first image of each subject was selected for the gallery, and the remaining images were used as probes. We achieved 100% rank-one recognition on this data set, even without masking fragile bits. Therefore, we cannot use rank-one-recognition-rate as a metric for showing the benefit of masking fragile bits. Instead, we looked at the separation between the genuine and impostor distributions of Hamming distance scores. One measure of separation between two distributions is the statistic d' (*d-prime*),¹²

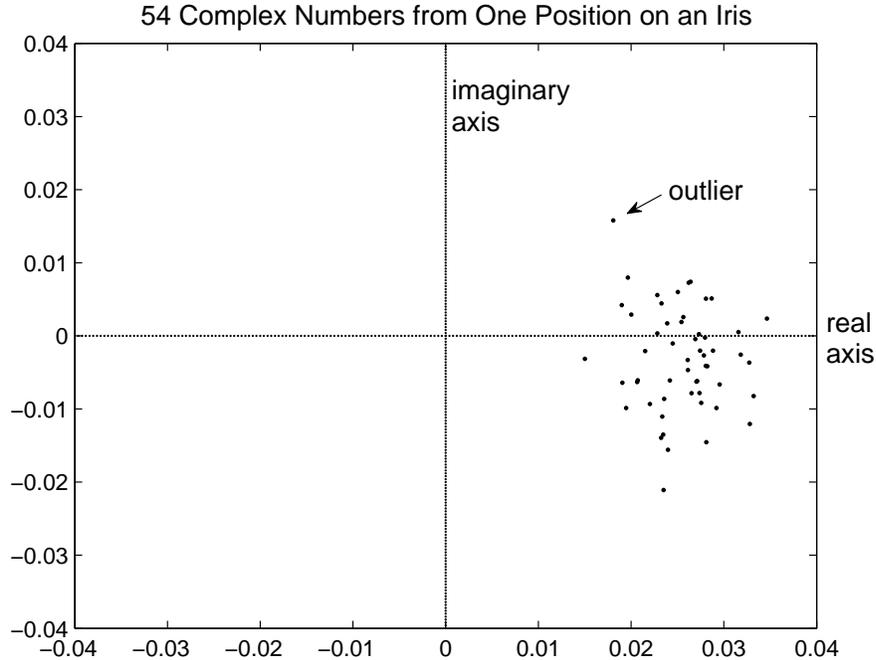


Figure 1. These 54 complex numbers, each from the same region on 54 different images of the same subject’s eye, all correspond to the same location on the iris code. Each complex number is mapped to two bits. This particular part of the iris code had a highly consistent real bit, and a highly inconsistent imaginary bit. (Figure reprinted from Hollingsworth et al., *IEEE Transactions on Pattern Analysis and Machine Intelligence*² ©2008 IEEE.)

which is defined as follows. If the means of the two distributions are μ_1 and μ_2 and their two standard deviations are σ_1 and σ_2 then d' is:

$$d' = \frac{|\mu_1 - \mu_2|}{\sqrt{\frac{1}{2}(\sigma_1^2 + \sigma_2^2)}}. \quad (1)$$

Since d' measures the separation between the genuine and impostor distributions, higher values are better.

We first ran an experiment using our original iris recognition algorithm on the 1226 images, and obtained a d' statistic of 7.48. Next, we implemented fragile-bit-masking, and ran an experiment using the same images, this time achieving a d' statistic of 8.25. Thus adding fragile-bit-masking to our algorithm increased the separation between the genuine and impostor distributions by more than three-fourths of the average standard deviation.

4. AVERAGE IRIS IMAGES FROM VIDEO

Although our algorithm described in the previous section works well, it is still limited by the quality of the images. Image quality is lower when an image is in poor focus, or when iris texture is hidden by specular highlights or eyelash occlusion. If a single still image is of poor quality, then we cannot have high confidence in the resulting verification or identification decision. With a video clip of an iris, however, a specular highlight in one frame may not be present in the next. Additionally, the amount of eyelash occlusion is not constant throughout all frames. It is possible to obtain a better image by using multiple frames from a video to create a single, less noisy iris image. Zhou and Chellappa¹³ suggested averaging to integrate texture information across multiple video frames to improve face recognition performance. We applied this averaging idea to iris recognition. Our proposed image-averaging technique is an example of signal-level fusion. A number of other researchers have investigated score-level fusion for iris recognition, but no prior work has attempted signal-level fusion.

In our proposed method,³ we selected N frames from a video clip of an iris. For each frame, we segmented and unwrapped the iris to create a normalized 20 by 240 representation of the texture in the iris. Then we averaged

together the unmasked pixels in the N normalized iris images. We applied this technique on 1061 videos from 296 different eyes, collected across a two-month time span, using $N = 6$. The first video of each eye was used as the gallery, and the remaining videos of each subject were used as probes.

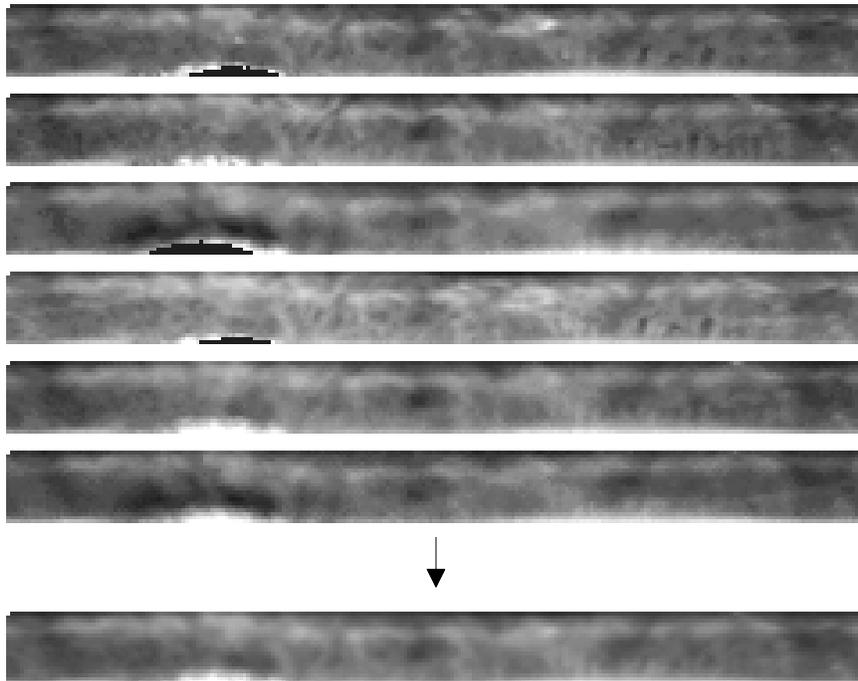


Figure 2. From the six original images on the top, we created the average image shown on the bottom.

We compared our method to score-level fusion approaches.^{14,15} We considered multi-gallery approaches that used N images for the gallery and a single image for the probe. Two ways of fusing Hamming distance scores are (1) taking the average of the N scores, or (2) taking the minimum of the N scores. Our method performed better than multi-gallery approaches using either an “average” or “minimum” score-fusion rule. Table I shows error rates for our proposed method in comparison with these other methods.

Table I: Signal-fusion Compared to Previous Methods

First Author	Method	EER	FRR @ FAR=0.001
Baseline	no fusion	1.64×10^{-2}	4.51×10^{-2}
Ma ¹⁴	score fusion: average	6.93×10^{-3}	1.40×10^{-2}
Krichen ¹⁵	score fusion: minimum	6.63×10^{-3}	1.43×10^{-2}
Proposed method	signal fusion	3.88×10^{-3}	7.61×10^{-3}

We conclude that on our data set, the signal-fusion method generally performs better than the previously proposed N-to-1-average or N-to-1-minimum methods. In addition, the signal fusion takes $1/N^{th}$ of the storage and $1/N^{th}$ of the matching time.

5. DETECTING IRIS TEXTURE DISTORTIONS FROM CONTACT LENSES

Image averaging mitigates the effects of poor quality iris images by averaging iris images from a single iris video. Another approach is to explicitly detect local regions of poor iris quality and drop those regions from our calculations. One example of a local texture distortion occurs with contact lenses that are manufactured with a

logo or symbol printed on them, such as the “AV” that is visible in the iris images in Figure 3a. This particular logo, the “AV,” is functionally important in that it is intended to help contact lens wearers to put the lens on correctly. Other local texture distortions can be caused by poorly-fitting contacts, the edges of hard contact lenses, segmentation inaccuracies, or shadows on the iris.⁴

Existing approaches to detecting local distortions in the iris texture focus on analysis of the iris image. In contrast, our approach⁴ to detecting local distortions focuses on analysis of the iris code matching results. Our approach has the advantage of making only the most general assumption about the cause of the local distortion in the iris texture. Also, this approach can be applied independently of and in combination with any improved iris segmentation algorithm.

Consider a comparison between two iris codes of the same iris. A local region in the comparison that has a high density of non-matching bits suggests that the corresponding region of one or both images contains a texture distortion or a segmentation inaccuracy. This suggests the following simple approach to detecting such regions:

1. Cover the match results with a large number of small windows.
2. Compute the fractional Hamming distance separately for each window.
3. Identify as “outlier” windows those that have an unusually high concentration of non-match bits.
4. Recompute the fractional Hamming distance without the contribution from the outlier windows.

This distortion detection will reduce the chances of a false reject.

We applied this approach to the iris images shown in Figure 3a. From each image, we segmented and unwrapped the iris to create a 20x240 representation of the iris texture; then we calculated the corresponding iris code. The fractional Hamming distance between these two images was 0.21.

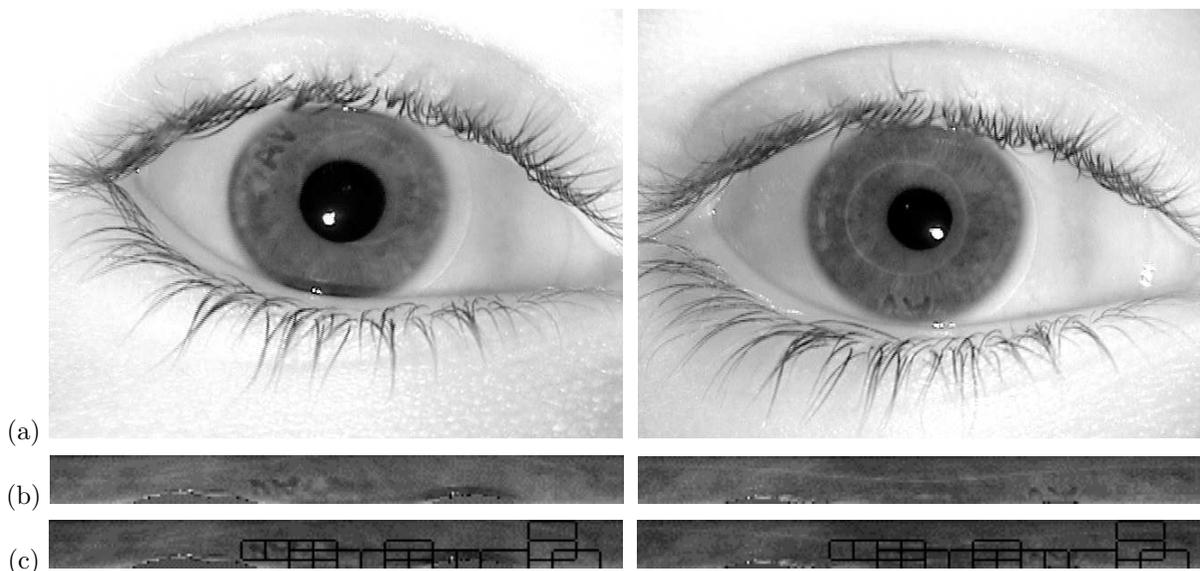


Figure 3. (a) Two images of the same iris (04855d115 and 04855d227), with the AV printed on the contact lens overlaying different parts of the iris. (b) The unwrapped iris texture for each image. (c) Locations of the twenty windows on each iris that have the highest fractional Hamming distance. The detected regions include the locations of the AV logo on each iris image, as well as the location of a shadow near the lower eyelid of the figure on the left. (Figures reprinted from Ring and Bowyer, *Proc. IEEE Conference on Biometrics: Theory, Applications, and Systems*⁴ ©2008 IEEE.)

To detect local texture distortions, we covered each 20x240 image with windows of size 8x20 that overlapped by one-half in each dimension. Thus there were $4 \times 23 = 92$ windows for each of the images. Next we computed the fractional Hamming distance for each window. Any window with fewer than 20 unmasked bit positions was dropped from consideration, because the statistics from such a small number of bits was found to be too noisy. The fractional Hamming distances for the remaining 8x20 windows ranged from a low of 0 to a high of nearly 0.5. Figure 3c shows the location of the twenty windows with the highest fractional Hamming distance. These

twenty windows include the location of the AV logo in each image, and also the location of a shadow near the lower eyelid of one of the images. Once these twenty windows were removed from the calculations, the fractional Hamming distance dropped from 0.21 to 0.15. If the chosen operating threshold of a system was anywhere in that range, our approach would prevent this user from being falsely rejected by the system.

As iris biometrics moves toward possible use in large-scale applications, it is important that no element of society be selectively disadvantaged in using the technology. It would be especially unfortunate if difficulty in using iris biometrics technology were correlated with a covariate such as the use of contact lenses. This method has the potential to reduce such problems or at least diagnose their manifestation.

6. EFFECTS OF PUPIL DILATION

Consider an application in which iris images are acquired under two different lighting environments. The system will acquire iris images of varying degrees of pupil dilation. Canonical iris biometrics algorithms account for dilation by “unwrapping” the iris image using Daugman’s rubber-sheet model.¹² Raw image coordinates are converted into normalized polar coordinates. Traditionally, all dilation information is discarded during this step.

The mapping to polar coordinates is important, because it makes possible a comparison between two different size irises, or between two images where the pupil has differing degrees of dilation. However, the rubber-sheet model assumes that iris tissue stretches linearly in the radial direction, an assumption which is not entirely accurate.

In order to study the effects of dilation on iris biometrics, we collected a data set of iris images with varying degrees of dilation.⁵ We acquired 1263 images from 18 different subjects (36 eyes). Some of the images were taken with the overhead room lighting turned off. The subjects’ pupils dilated because of the lack of visible light, but the irises were still illuminated using infrared LEDs. We measured the dilation ratio for each image. We defined dilation ratio to be the pupil radius divided by the iris radius. All dilation ratios in our dataset fell between 0.2 and 0.8. Some sample images are shown in Figure 4.

We divided our data into three subsets: one subset with small pupils, one with medium pupils, and one with large pupils. After running the experiment, we found that the subset of our data with all large pupils showed the worst performance. The mean of the distribution of genuine Hamming distance scores was 0.06 higher, and the distribution of impostor Hamming distance scores was 0.02 lower, when compared to the subset of the data with small pupils. The equal error rate was an order of magnitude greater for the subset of data with large pupils, compared to the small-pupil data set. The decision error threshold curves are shown in Figure 5.



Figure 4. This subject showed the biggest difference in pupil size in the data set. The smallest dilation ratio (pupil radius/iris radius) for this subject was 0.3478 and the largest dilation ratio was 0.6545. (Figure reprinted from Hollingsworth et al., *Proc. IEEE Conference on Biometrics: Theory, Applications, and Systems*¹⁶ ©2008 IEEE.)

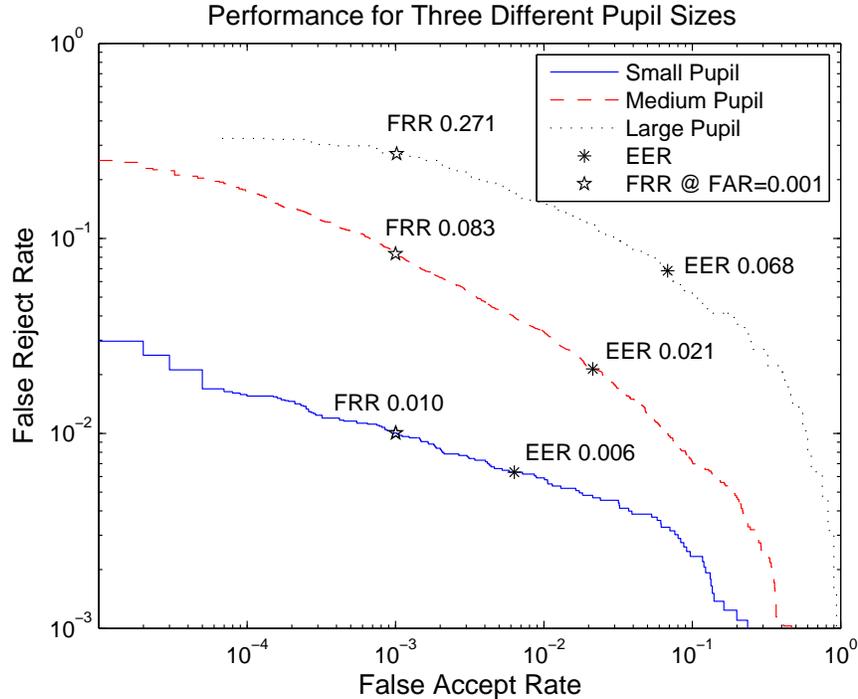


Figure 5. Our iris biometric algorithm performs significantly better when the data set contains iris images with all small pupils. (Figure reprinted from Hollingsworth et al., *Computer Vision and Image Understanding*,⁵ with permission from Elsevier.)

We considered two factors to explain this result. First, when pupils are dilated, there is less iris area visible. With fewer pixels in the iris region, there is less information to accurately characterize the texture of the iris. Second, when pupils are dilated, a greater part of the iris is pulled towards the eyelid. Thus, a larger percentage of iris area is occluded. The percentage of occluded area is important because the impostor distribution is affected by the amount of unmasked data available. If the system bases a comparison on fewer bits, the impostor distribution will have a greater standard deviation, and thus more false accepts. This phenomenon encourages the idea of score normalization as proposed by Daugman.¹⁷ With score normalization, the fractional Hamming distance is normalized by the number of unmasked bits available for comparison. This score normalization reduces false accepts.

In our second experiment, we asked the question, “do two images with the same dilation match better than two images with different dilation?” For this experiment, we considered all comparisons possible in our data. We divided the comparisons based on the difference of dilation ratio between the two eyes in the comparison. We found that genuine comparisons from images with the same degree of dilation had smaller Hamming distances than genuine comparisons between images with widely different degrees of dilation. The Hamming distance distribution for comparisons with large differences in dilation had a mean about 0.08 higher than the distribution involving comparisons with the same degrees of dilation. The distribution of impostor Hamming distances was unaffected. The equal error rate for the comparisons with different degrees of dilation was nearly four times the equal error rate for the comparisons with the same degree of dilation.

We conclude that iris biometrics performs significantly worse on data sets containing significant pupil dilation. The order of magnitude difference in equal error rate from our first experiment, or the four times difference in equal error rate from our second experiment cannot be easily ignored. We recommend that a measure of pupil dilation for each iris should be incorporated into a quality metric for the match. Score normalization should be used to help prevent false accepts.

To improve performance, we could purposely enroll multiple images of a subject with varying degrees of

dilation. We could also look at other ways to account for dilation. Thornton et al.¹⁸ and Wei et al.¹⁹ have already begun looking at ways to model the deformations caused by dilation.

7. EFFECTS OF TIME LAPSE

Aside from pupil dilation and barring any traumatic injury to the eye or interocular surgery, the iris is assumed to remain stable over time. However, to our knowledge, there has been no study of the long-term stability of iris texture as imaged under NIR illumination.

We use a data set with approximately four years time lapse for 23 subjects, 46 iris-subjects, to test the stability of the iris over time. For each iris-subject we compare two types of matches: (1) matches between two images acquired within 120 days of each other and (2) matches between images acquired with more than 1200 days time-lapse. Using these sets of long-time-lapse matches and short-time-lapse matches, we test the null hypothesis that the normalized Hamming distance for long-time-lapse matches is not different from that for short-time-lapse matches.

We considered two experimental scenarios to test the null hypothesis, an “all-irises” test and an “iris-level” test. For the “all-irises” scenario, we found each iris-subject’s average Hamming distance for short-time-lapse matches, μ_{S*S} , and the average Hamming distance for long-time-lapse matches, μ_{S*L} . We computed the difference between the means, $\mu_{S*L} - \mu_{S*S}$, for each iris-subject, yielding a set of 46 differences of means. A histogram representing this sample of differences of mean Hamming distances is shown in Figure 6. The difference of means was positive for 43 of the 46 irises.

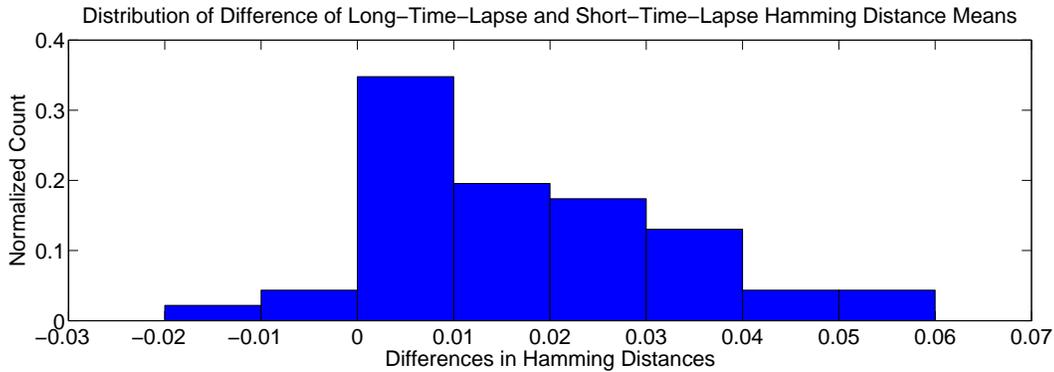


Figure 6. This distribution of difference of long-time-lapse and short-time-lapse is clearly shifted to the right of zero.

In order to test the null hypothesis for the “all-irises” scenario, we applied a sign test to test the null hypothesis that the number of positive differences of means is not statistically greater than the number of negative differences. We also applied a Student’s t test to test the null hypothesis that the difference-of-means sample comes from a distribution with a mean of zero. Both statistical tests rejected the null hypothesis at a 5% significance level. We concluded that the sample of long-time-lapse mean Hamming distances came from a distribution with a mean greater than the sample of short-time-lapse mean Hamming distances.

In the second scenario, the “iris-level” test, we used the sample of long-time-lapse Hamming distances and the sample of short-time-lapse Hamming distances for each iris-subject. Applying a one-tailed Student’s t test, we test the null hypothesis that these two samples of Hamming distances come from the same distribution with equal means against the alternative hypothesis that the sample of of long-time-lapse matches is greater than the short-time-lapse matches. The null hypothesis was rejected for 40 of the 46 irises at a significance level of 5%.

We consider several factors other than time-lapse that could conceivably affect our results. Throughout our experiments, we implemented score normalization¹⁷ and fragile bit masking to control for the number and fragility of bits used in a match comparison. We determined that the pupil to iris ratio difference between images in a match did not affect the degrading match quality in long-time-lapse matches. We checked for any change

in contact lenses for each subject over the four year period. We performed experiments to test the effect of the camera position in our acquisition studio and the potential aging of the sensor on match quality. We conclude that evidence exists to suggest iris match quality degrades with increased time-lapse between image acquisitions. The exact cause of this observed degradation is not yet clear and the amount of degradation in the average Hamming distance is not large at the four-year mark. To account for the time-lapse effect, we could re-enroll a subject with every verification scenario or require re-enrollment for each subject after a set time frame.

8. CONCLUSION

We recommended a number of ways to improve iris biometric performance. The use of any one of these recommendations does not preclude the use of any of the others. We presented a method to mask out fragile bits in the iris code. This method eliminates the effects of the inconsistencies in the iris code that arise from the quantization of the complex filter response in a canonical iris biometrics algorithm. Second, we showed that signal-level fusion of multiple frames in an iris video provides better performance than using single still iris images. The fusion of multiple frames removes much of the noise in the iris images while retaining relevant texture information. Third, local problem areas in iris images can be detected and removed from computations. To detect local distortions, we broke unwrapped iris images into small windows and computed Hamming distances for each window; windows with high concentrations of non-matching bits were removed.

In our study involving images taken under varying lighting conditions, we found that images of dilated iris do not work as well for iris biometrics. We recommended using Daugman's method of score normalization to reduce false accepts. Improving recognition for dilated irises remains an open area of study.

Our final study involved matching iris images acquired four years apart. We found that a long time lapse between acquisitions harmed performance. The exact cause of this effect is another open area of study.

ACKNOWLEDGMENTS

This research is supported by the National Science Foundation under grant CNS01-30839, by the Central Intelligence Agency, by the Intelligence Advanced Research Projects Activity, by the Biometrics Task Force, and by the Technical Support Working Group under US Army contract W91CRB-08-C-0093. The opinions, findings, and conclusions or recommendations expressed in this publication are those of the authors and do not necessarily reflect the views of our sponsors.

REFERENCES

- [1] Bowyer, K. W., Hollingsworth, K. P., and Flynn, P. J., "Image understanding for iris biometrics: A survey," *Computer Vision and Image Understanding* **110**(2), 281–307 (2008).
- [2] Hollingsworth, K. P., Bowyer, K. W., and Flynn, P. J., "The best bits in an iris code," *IEEE Transactions on Pattern Analysis and Machine Intelligence*. accepted for publication.
- [3] Hollingsworth, K. P., Bowyer, K. W., and Flynn, P. J., "Image averaging for improved iris recognition," in *Proc. Int. Conf. on Biometrics (ICB2009)*, (2009).
- [4] Ring, S. and Bowyer, K. W., "Detection of iris texture distortions by analyzing iris code matching results," in *Proc. IEEE Int. Conf. on Biometrics: Theory, Applications, and Systems*, (Sept 2008).
- [5] Hollingsworth, K. P., Bowyer, K. W., and Flynn, P. J., "Pupil dilation degrades iris biometric performance," *Computer Vision and Image Understanding* **113**(1), 150–157 (2009).
- [6] Baker, S., Bowyer, K. W., and Flynn, P. J., "Empirical evidence for correct iris match score degradation with increased time-lapse between gallery and probe matches," in *Proc. Int. Conf. on Biometrics (ICB2009)*, (2009).
- [7] Phillips, P. J., Scruggs, W. T., O'Toole, A. J., Flynn, P. J., Bowyer, K., Schott, C. L., and Sharpe, M., "FRVT 2006 and ICE 2006 large-scale results," tech. rep., National Institute of Standards and Technology, NISTIR 7408 (Mar 2007). <http://iris.nist.gov/ice>.
- [8] Matey, J. R., Naroditsky, O., Hanna, K., Kolczynski, R., LoIacono, D., Mangru, S., Tinker, M., Zappia, T., and Zhao, W. Y., "Iris on the MoveTM: Acquisition of images for iris recognition in less constrained environments," *Proceedings of the IEEE* **94**(11), 1936–1946 (2006).

- [9] Phillips, P. J., Bowyer, K. W., Flynn, P. J., Liu, X., and Scruggs, W. T., “The iris challenge evaluation 2005,” in *Proc. IEEE Int. Conf. on Biometrics: Theory, Applications, and Systems*, (Sept 2008).
- [10] Phillips, P. J., Scruggs, T., Flynn, P. J., Bowyer, K. W., Beveridge, R., Givens, G., Draper, B., and O’Toole, A., “Overview of the multiple biometric grand challenge,” in *Proc. Int. Conf. on Biometrics (ICB2009)*, (2009).
- [11] Liu, X., Bowyer, K. W., and Flynn, P. J., “Experiments with an improved iris segmentation algorithm,” in *Proc. Fourth IEEE Workshop on Automatic Identification Technologies*, 118–123 (Oct 2005).
- [12] Daugman, J., “How iris recognition works,” *IEEE Transactions on Circuits and Systems for Video Technology* **14**(1), 21–30 (2004).
- [13] Zhao, W. and Chellappa, R., eds., [*Face Processing: Advanced Modeling and Methods*], ch. 17: Beyond one still image: Face recognition from multiple still images or a video sequence by S.K. Zhou and R. Chellappa, 547–567, Elsevier (2006).
- [14] Ma, L., Tan, T., Wang, Y., and Zhang, D., “Efficient iris recognition by characterizing key local variations,” *IEEE Transactions on Image Processing* **13**, 739–750 (Jun 2004).
- [15] Krichen, E., Allano, L., Garcia-Salicetti, S., and Dorizzi, B., “Specific texture analysis for iris recognition,” in *Int. Conf. on Audio- and Video-Based Biometric Person Authentication (AVBPA 2005)*, 23–30 (2005).
- [16] Hollingsworth, K. P., Bowyer, K. W., and Flynn, P. J., “The importance of small pupils: a study of how pupil dilation affects iris biometrics,” in *Proc. IEEE Int. Conf. on Biometrics: Theory, Applications, and Systems*, (Sept 2008).
- [17] Daugman, J., “New methods in iris recognition,” *IEEE Transactions on Systems, Man and Cybernetics - B* **37**, 1167–1175 (Oct 2007).
- [18] Thornton, J., Savvides, M., and Kumar, B. V., “A Bayesian approach to deformed pattern matching of images,” *IEEE Transactions on Pattern Analysis and Machine Intelligence* **29**, 596–606 (Apr 2007).
- [19] Wei, Z., Tan, T., and Sun, Z., “Nonlinear iris deformation correction based on Gaussian model,” in *Proc. Int. Conf. on Biometrics (ICB2007)*, 780–789 (Aug 2007).