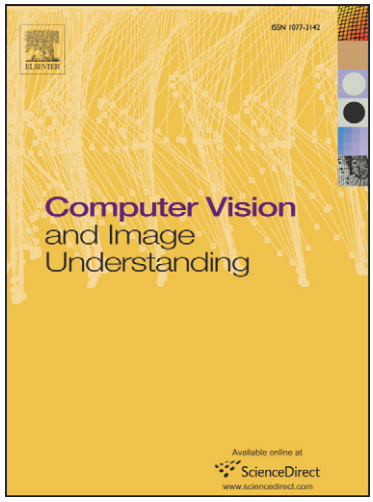

Pupil Dilation Degrades Iris Biometric Performance

Karen Hollingsworth, Kevin W. Bowyer, Patrick J. Flynn

PII: S1077-3142(08)00117-3
DOI: 10.1016/j.cviu.2008.08.001
Reference: YCVIU 1473

To appear in: *Computer Vision and Image Understanding*

Received Date: 14 January 2008
Revised Date: 9 June 2008
Accepted Date: 4 August 2008

Please cite this article as: K. Hollingsworth, K.W. Bowyer, P.J. Flynn, Pupil Dilation Degrades Iris Biometric Performance, *Computer Vision and Image Understanding* (2008), doi: 10.1016/j.cviu.2008.08.001

Pupil Dilation Degrades Iris Biometric Performance

Karen Hollingsworth, Kevin W. Bowyer, and Patrick J. Flynn

Dept. of Computer Science and Engineering,

University of Notre Dame

Notre Dame, Indiana 46556

kholling, kwb, or flynn @ cse.nd.edu

June 9, 2008

Abstract

Iris biometrics research has largely ignored the problems associated with variations in pupil dilation between the enrollment image and the image to be recognized or verified. Indeed, in most current systems, information about pupil dilation is discarded when the iris region is normalized to a dimensionless polar coordinate system from which the iris code is obtained. This work studies the effect of pupil dilation on the accuracy of iris biometrics. We found that when the degree of dilation is similar at enrollment and recognition, comparisons involving highly dilated pupils result in worse recognition performance than comparisons involving constricted pupils. We also found that when the matched images have similarly highly dilated pupils, the mean Hamming distance of the match distribution increases and the mean Hamming distance of the non-match distribution decreases, bringing the distributions closer together from both directions. We further found that when matching enrollment and recognition images of the same person, larger differences in pupil dilation yield higher template dissimilarities, and so a greater chance of a false non-match. We recommend that a measure of pupil dilation be kept as meta-data for every iris code. Also, the absolute dilation of the two images, and the dilation difference between them, should factor into a confidence measure for an iris match.

Keywords: Iris biometrics, pupil dilation.

1 Introduction

Each individual human iris is believed to have a unique texture pattern of sufficient complexity that it can be used for identification. Several commercial iris biometric systems exist [1, 2, 3, 4]. For a recent comprehensive survey of commercial and research systems for iris recognition, see [5]. Most commercial iris technology is based on algorithms developed by Daugman [6, 7]. These algorithms take an image of an eye, locate the limbic and pupillary boundaries of the iris, and apply a transformation that maps the iris into a dimensionless polar coordinate system. Next, Gabor texture filters are applied to generate a binary iris code that represents the texture in the iris. An iris code obtained from an image can be compared to a reference iris code for a person to verify a claimed identity, or compared to a gallery of iris codes to recognize a person.

One difficulty in processing iris images for biometrics is that the iris changes in size due to involuntary dilation. Two muscle systems, a sphincter and several radial dilator muscles, control the size of the iris to adjust the amount of light entering the pupil. Naturally, a recognition system should be made to be robust to changes in pupil dilation. In Daugman's early work [6], he assumed a "homogenous rubber sheet model" which assigns to each point on the iris a pair of coordinates (r, θ). The radius, r, ranges from 0 to 1, where points on the pupillary boundary have a radial coordinate of 0, and points on the limbic boundary have a radial coordinate of 1. The angle θ ranges from 0 to 2π. For an image $I(x, y)$, the mapping of raw coordinates (x, y) to dimensionless polar coordinates (r, θ) can be represented as $I(x(r, \theta), y(r, \theta)) \rightarrow I(r, \theta)$. This mapping results in a size-normalized representation of the iris. By interpreting the radius and angle as Cartesian coordiates, we can display a rendering of the size-normalized iris region. An example iris image from our dataset, and its corresponding size-normalized iris region are shown in Figures 1 and 2, respectively.[1]

The normalization of the iris makes it possible to compare two images of different sizes, but it discards information about the degree of pupil dilation. Furthermore, this transformation is not entirely accurate because it assumes that when the pupil dilates, the stretch of the iris tissue in the radial direction is linear. Some researchers have investigated other ways to account for pupil dilation more precisely, but no prior research has quantified how differing degrees of dilation affect the performance of an iris biometrics system. In this paper, we quantify the effects of pupil dilation on the authentic and imposter distributions in iris biometrics.

[1] In order to render this polar coordinate system on paper, we have introduced an arbitrary cut in the iris region between 0 and 2π. However, in actuality, there is no such cut in the angular variable.

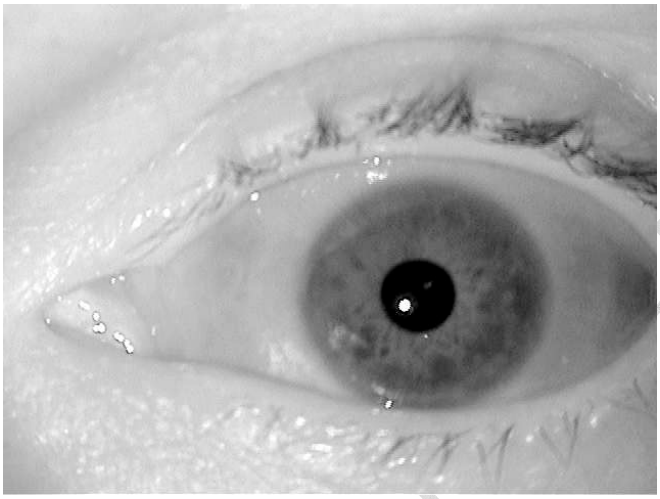

Figure 1: A sample iris image from our dataset.

Figure 2: A size-normalized rendering of the iris region from Figure 1.

2 Research in Iris Dilation

Iris tissue does not follow a perfect "rubber sheet" model when the pupil dilates and contracts. Wyatt [8] explains that this assumption is a good approximation, but does not perfectly match the actual deformation. He developed a mathematical model to explain how the collagen fibers in the iris deformed as the pupil dilated. Initially he restricted his model to require linear deformation along the radial direction of the iris, and then he later relaxed that constraint. He compared his models to measurements of surface features from several human irises. He reports that "some of the data ... appear to be a better match to the lines indicating linear behavior; other data ... appear to be a better match to the lines indicating nonlinear behavior." He further compared his model to some measurements of angles of collagen fibers measured on an iris, and found that the nonlinear model matched better.

Some iris biometrics researchers have noted that pupil dilation affects the quality of a match between iris images [9, 10, 11]. Ma et al. [9] characterized how many of their false non-matches were due to pupil dilation. They explained that "under the extreme conditions (namely the iris texture is excessively compressed by the pupil), the iris after normalization still has many differences with its normal state (i.e., the iris has a pupil of normal size). Therefore, the matching distance between such a pair of iris images is very large. In our experiments, 10.7% false non-matches result from the pupil changes. This is a common problem in all iris recognition methods." Thornton et al. [10] add an extra processing step to account for the nonlinear deformations of the iris that occur when the pupil dilates. They find the maximum a posteriori probability estimate of the parameters of the relative deformation between a pair of images. Their results show that estimating the relative deformation between the two images improves performance. Wei et al. [11] account for dilation by modeling nonlinear iris stretch as a sum of linear stretch and a Gaussian deviation term. Their training set includes multiple images of a subject taken under gradually varying illumination. They compare their algorithm with two previous algorithms which use the simple rubber-sheet model and show that their model achieves a lower equal error rate. Both Thornton and Wei focus on comparing algorithms. Neither work reports experimental results employing subsets of images with different degrees of dilation.

Other than the few papers mentioned above [9, 10, 11], the large majority of iris biometrics literature assumes that the "rubber sheet" approach of Daugman is sufficient to deal with the differences in dilation. Furthermore, even though a few researchers have looked at pupil dilation, we have not found any work that quantifies the effect of dilation on recognition. This work is the first that we are aware of to examine the impact of dilation on the performance of an iris recognition algorithm.

3 Data and Software

In order to measure how dilation affects recognition, a special iris image data set was collected at the University of Notre Dame between July 2007 and September 2007, using an LG 2200 EOU camera. This data set includes 630 left eye images and 633 right eye images from 18 different subjects. This set includes ten males and eight females. Twelve subjects were Caucasians and six were Asians. In order to have images of varying pupil dilation, we turned off the ambient lighting when acquiring 28% of the images. The LG 2200, like all commercial iris cameras, uses infrared LEDs to actively illuminate the iris and therefore can still take high-quality iris images in the absence of ambient lighting. Iris dilation is driven by the visible light level, not by infrared light. The LEDs in the LG 2200, while visible, are not bright enough to significantly affect pupil size.

We used iris recognition software developed by Liu [12] and modified as described in [13] and [14]. This software, based on ideas from Daugman [6] and Wildes [15] and on the implementation of Masek [16], uses a Canny edge detector and a Hough transform for segmentation. The boundaries of the iris are approximated using two circles that are not necessarily concentric. We visually inspected the segmentation for all the 1263 images. The iris and pupil were correctly located in all of the images. Once an image is segmented, log-Gabor filters are used to analyze the texture of the iris and compute a binary iris code. If parts of the iris are occluded by eyelids or eyelashes, the corresponding bits in the iris code are masked, or excluded from future computations. To compare two irises, our system calculates the fractional Hamming distance, or the fraction of unmasked bits that disagree between the two iris codes. For conciseness throughout this paper, any reference to Hamming distance refers to the fractional Hamming distance.

To show how this software performs for iris recognition, a graph of the match and non-match distributions for this image set is shown in Figure 3. The match distribution shows the histogram of Hamming distances between iris codes for comparisons in which both images are of the same iris. The non-match distribution shows distances from comparisons in which the two images come from different irises. The amount of overlap between the two distributions is related to the error rates in the system. As the figure shows, the match distribution clearly has a lower mean than the non-match distribution, but there is some overlap between the distributions. In commercial biometric systems, this overlap is typically handled by setting a threshold on the Hamming distance so that the probability of a false match is at or below a specified level; e.g., less than one in one million.

Daugman [7] documents the expected non-match distribution for an iris recognition system. He explains

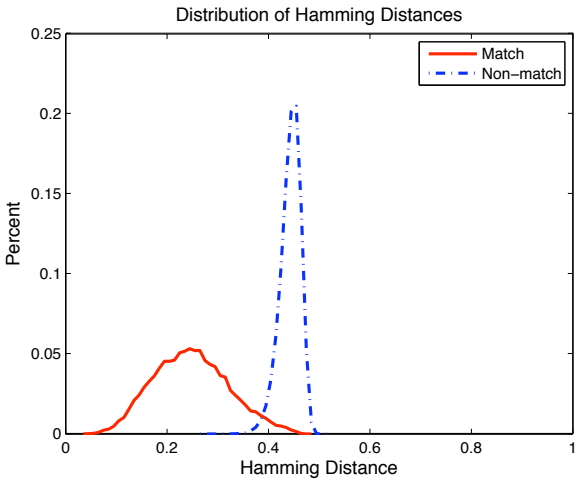

Figure 3: Distribution of fractional Hamming distances for all iris comparisons in our data set. Two histograms are shown. The one on the left represents comparisons between two images of the same iris. The one on the right represents comparisons between two different irises.

that "because any given bit in the phase code for an iris is equally likely to be a 1 or 0, and different irises are uncorrelated, the expected proportion of agreeing bits between the codes for two different irises is HD = 0.500." However, in order to account for tilt of a person's head when an image is taken, many iris recognition algorithms try multiple possible rotations of an iris during a comparison, and assume the best possible match to be the correct alignment of the iris. Daugman notes that "this 'best of n' test skews the distribution to the left and reduces its mean from about 0.5 to 0.458" [7]. Likewise, our non-match distribution is also skewed slightly to the left, with a mean of 0.443, and with the majority of Hamming distance values falling between 0.4 and 0.5.

Our match distribution shows that our software does not perform as well as Daugman's. In Figure 9 of Daugman's paper [7], his match distribution has a mean Hamming distance of 0.110. The mean Hamming distance for our match distribution is 0.2479. However, our dataset intentionally includes wide variation in pupil dilation. We argue that our software performs well enough that it can be used to show interesting observations about how dilation affects the performance. As demonstrated in [12], this system achieves a 97.1% rank-one recognition rate on an experiment of 4249 images.

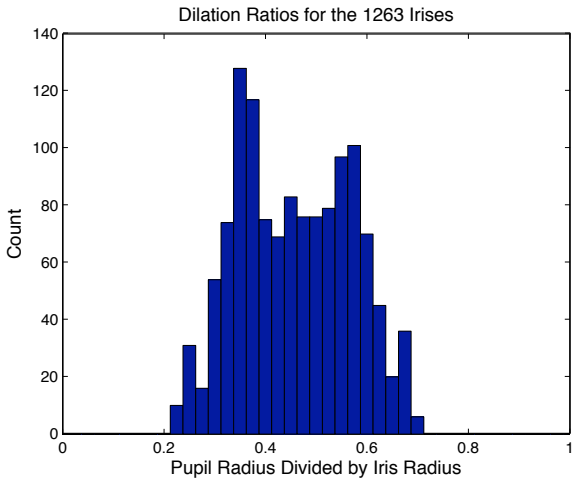

Figure 4: Histogram showing different degrees of dilation for irises in our data set. A dilation ratio of 0.2 corresponds to a very small pupil (contracted eye), and a dilation ratio of 0.7 corresponds to a very large pupil (dilated eye).

4 Measuring Dilation

Our first step was to measure the degree of dilation for each image. The segmentation step of the analysis provided the radius of the pupil and the radius of the iris. To measure dilation, we divided the pupil radius by the iris radius. Since the pupil radius is always less than the iris radius, this dilation ratio must fall between 0 and 1. In our 1263 images, all dilation ratios were between 0.2137 and 0.7009. The distribution of dilation ratios is shown in Figure 4.

Winn et al. [17] report that "several factors are known to affect pupil size, including the level of retinal illuminance, the accommodative state of the eye, and various sensory and emotional conditions. In addition, the size of the pupil tends to change as a function of the individual's age, with smaller pupils being predominant in the elderly population." In our data, the two oldest subjects had the smallest pupil size. The minimum dilation ratio for the oldest subject was 0.21. Wyatt [8] reports that typical pupil diameters during waking hours fall within 12 to 60 percent of iris diameters. We do not have any images with a pupil diameter as small as 12% of the iris diameter. However, since pupil size is somewhat correlated with age [17], the lack of images with very small pupils may be due to the fact that the majority of the subjects in our dataset are between 19 and 32 years old, and we have no subjects older than 52 in our dataset.

Two of the subjects attended only one acquisition session. Figure 5 plots the minimum and maximum

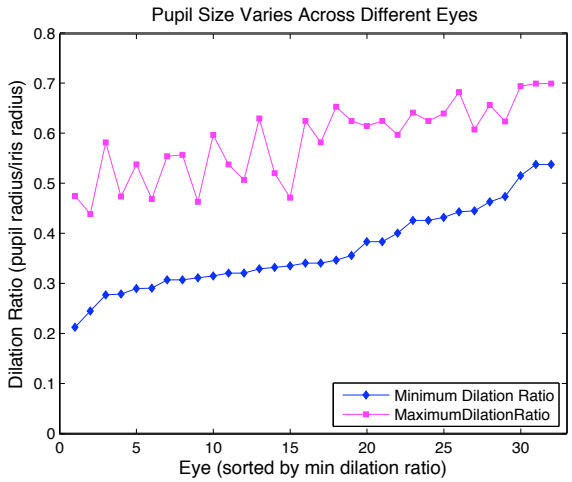

Figure 5: Minimum and maximum dilation ratios for 32 different eyes, sorted by minimum dilation ratio. Some subjects' eyes naturally tended to have higher dilation ratios than other subjects' eyes. Also, some subjects showed more variation in their pupil size than others.

dilation ratios for the remaining 16 subjects (32 eyes). For one eye, the pupil radius varied as much as 31% of the iris radius. The eye showing the least change in pupil dilation still had the pupil radius varying by 14%.

As mentioned in the previous section, some of the images were taken with the overhead lights in the room off. Therefore it is natural to wonder whether the variation in pupil size might occur under real-world conditions. One realistic scenario that causes pupils to contract is a higher level of illumination in the room. One simple way to get pupils to dilate is to use eye drops containing atropine sulphate, a chemical compound used by optometrists to dilate the pupil during eye exams. Such drops can cause a person's pupils to dilate and remain dilated for hours. An attacker might use such drops to try to fool the biometric system, or a legitimate user may have dilated pupils after a routine visit to the optometrist. An even more common way of causing pupils to dilate is to wear sunglasses. Sunglasses typically block harmful ultraviolet frequencies of light, but have no effect on infrared illumination. Therefore, an infrared image of the eye can be acquired even when a subject is wearing sunglasses. We took two pictures, one of an author wearing sunglasses and one of the same person with an extra lamp next to the camera (Figure 6). With sunglasses, the pupil dilation ratio was 0.54 which was nearly as large as the largest dilation ratio from that subject with the lights off (0.56). With the extra lamp in the room, the dilation ratio was 0.31. The pupil radius had varied by 23% between the two pictures. Therefore, we conclude that such a range of pupil variation is not uncommon and

should be expected and evaluated.

We avoided using sunglasses in acquiring the images for our experiments, because we did not want the complicating factor of extra specular reflections from the glasses. Therefore, all the images used in our experiments were taken without glasses. Any differences in dilation were obtained by turning off the ambient lighting. Figures 7, 8, and 9 show the pupil variation for three of the subjects in our dataset. These three figures show two iris images each: one image taken under normal ambient room lighting, and one taken with the ambient lights off.

5 Degree of Dilation Affects Performance

In evaluating the impact of dilation on iris biometrics, the first question to consider is whether the iris recognition software performs well on a set of data made up of entirely dilated pupils. One might expect that if all images in a dataset showed a large degree of pupil dilation, but were all consistently dilated, the system would still perform well.

Since all dilation ratios in our dataset fall between 0.2 and 0.8, we divided the images into three sets. Set 1 contains images with a dilation ratio less than 0.4. Images in this set have small pupils. Set 2 contains images with a dilation ratio between 0.4 and 0.6. Finally, set 3 contains the most dilated pupils, with dilation ratios greater than 0.6. We obtained the match and non-match distributions for each of the three sets. The match distributions for all three sets are shown in Figure 10.

Surprisingly, even though all irises within each set have consistent degrees of dilation, not all sets showed the same performance. In set 1, the mean Hamming distance for the match distribution is 0.1982. For set 2, the mean Hamming distance is 0.2470. Set 3 has a mean Hamming distance of 0.2592. It is possible that this result (and other results presented in this paper) may be due in part to confounding factors such as focus, or imperfections in eyelash detection. However, there are images with varying degrees of focus and varying amounts of eyelash occlusion in all three sets; thus we do not suspect that these would be the main factors in this result. We conclude then, that as pupils get larger, the mean of the match distribution increases, getting closer to the non-match distribution and increasing the rate of false rejects, or false non-matches.

One probable reason for the degraded performance in the set of dilated eyes involves the simple fact that there is less iris area visible. Typical iris recognition algorithms convert the annular region of the eye to a size-normalized region. In a dilated image, less iris area is available for creating each pixel in the size-normalized image. With less iris data available the eyes will be characterized more poorly.

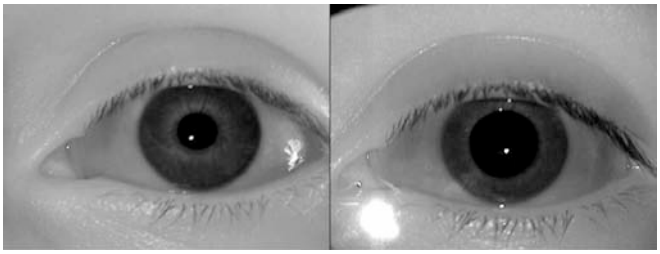

(a) (b)

Figure 6: It is easy to obtain a difference in pupil size. The iris image in part (a) was taken with an extra lamp in the room, and the image in (b) was taken with the subject wearing sunglasses. The pupil in (b) is almost as dilated (dilation ratio 0.54) as the most dilated pupil (0.56) seen in our experimental data for this subject.

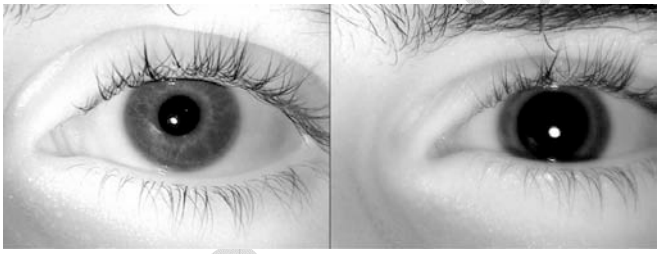

(a) (b)

Figure 7: This subject (subject number 05288) 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.

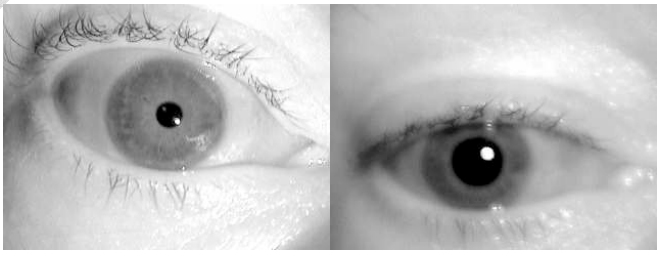

(a) (b)

Figure 8: This subject (subject number 02463) had the smallest pupils in the dataset. The smallest dilation ratio for this subject was 0.2137, and the largest dilation ratio was 0.4762.

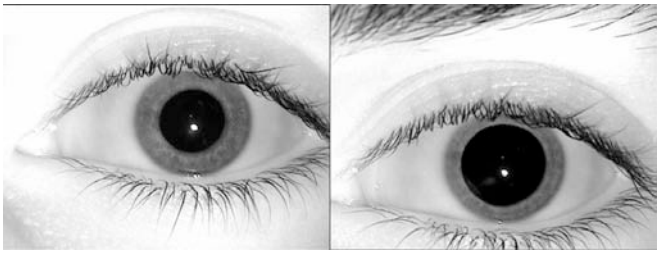

(a) (b)

Figure 9: This subject (subject number 05379) had the largest pupils in the dataset. The smallest dilation ratio for this subject was 0.5391, and the largest was 0.7009.

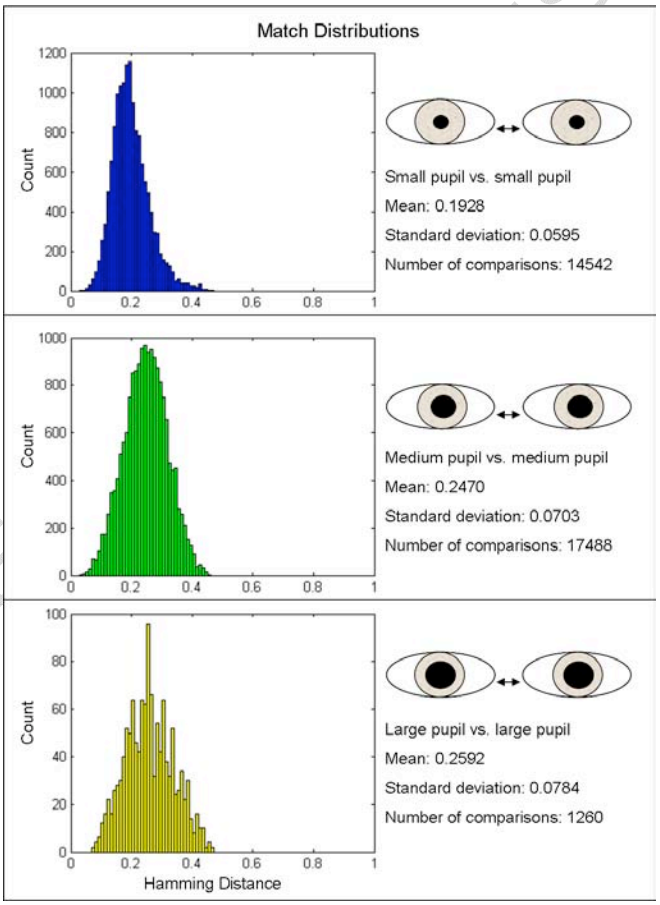

Figure 10: The match distribution for comparisons between eyes with large pupils has a larger mean than the distribution of comparisons between eyes with small pupils.

The creation of the size-normalized iris image involves sampling the original image. The sample points on the original image will be spaced out along a line going from the pupil, radially outward. A typical iris image might have sample points approximately three pixels apart. For an eye with a large pupil, the sampling density along such a line will be higher. As pupil dilation increases, the radial width of the iris decreases. Thus, there are fewer distinct pixels along the annular width of the iris, but the number of sample points for the normalized iris image remains the same.

In creating the normalized image, our system uses a square of four pixels and interpolates a single value from those four pixels to create one pixel of the normalized image. The normalized image that is created is 20 pixels by 240 pixels. Therefore, the system is expecting an iris with a width at least 40 pixels across. The most dilated pupil in the images analyzed for this paper has an iris that is only 35 pixels in annular width. For this image, the values of some pixels are sampled more than once in creating the normalized image.

The International Standards Organization has specified that an iris image used for recognition should have 200 pixels across the diameter of the iris [18]. However, we suggest that the diameter of the iris is an inadequate way to measure the total amount of iris data available. Instead, we recommend that the annular width of the iris is a more correct measure of iris size. Annular width is easily computed by taking the iris radius and subtracting the pupil radius. For systems that allow non-circular pupillary and/or limbic boundaries, the distance between the two can be averaged over a number of radial samples. Even if two circles describing the pupil and iris are not concentric, this measure of annular width is still easily computed, and still represents an average annular width of the iris. However, recording the minimum annular width may also be useful in determining the quality of an image.

Figure 11 shows the non-match distributions for the three sets of iris images. In set 1, the mean Hamming distance for the non-match distribution is 0.4483. The mean Hamming distance for set 2 is 0.4398 and the mean Hamming distance for set 3 is 0.4265. The difference in means for the non-match distribution was not as large as the difference in means for the match distribution. Therefore, we wished to test whether this difference was statistically significant. To simplify the test, we split the data into two groups, where the first group contains all eyes with dilation ratios below the median, and the second group contains all eyes with dilation ratios above the median. In this test, the null hypothesis is that the means of the non-match distributions for the two groups are equal, and the alternative hypothesis is that they are not equal. We used a balanced, one-factor ANOVA test to compare the means of the two groups of data. The factor in this test is the dilation ratio. We got an F-statistic of 51512, and a corresponding p-value of 0.0000. Thus, we concluded that the difference between the two groups was indeed significant. That is, the non-match

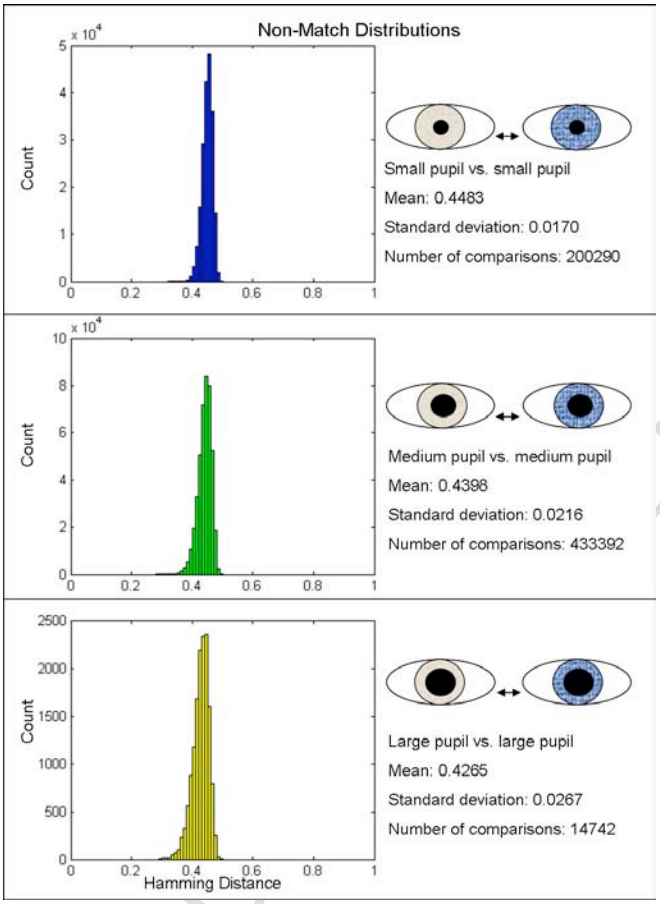

Figure 11: The non-match distribution of comparisons between eyes with large pupils has a smaller mean than the distribution of comparisons between eyes with small pupils.

distribution for dilated eyes has a different mean than the non-match distribution for non-dilated eyes. In our experimental results, the mean Hamming distance of the non-match distribution decreases as pupils get larger, moving closer to the match distribution and therefore increasing the rate of false matches.

We know of two possible explanations for the decrease in the mean of the non-match distributions. The first explanation relates to the number of bits used in a comparison between two iris codes. Our system employs masking logic to exclude parts of the iris code that are affected by eyelids and eyelashes. This masking leaves fewer bits available for the Hamming distance computation when an iris has more eyelid and eyelash occlusion. Pupil dilation pulls more of the iris towards the eyelids, causing a larger percentage of normalized iris area to be occluded, and consequently fewer bits used in the comparison.

When fewer bits are available for comparison between two iris codes, one expects the underlying raw distribution of fractional Hamming distance scores (before rotations to allow for head tilt) to be broader

than when more bits are unmasked and available. This is for the same reason that tossing a coin many times will generate a narrower distribution of outcome fractions of "heads" than when tossing it few times; indeed the standard deviation of that fractional binomial distribution in [0,1] is inversely proportional to square root of n, where n is the number of coin tosses, or the number of bits. Therefore, large pupil dilation, which reduces n, increases the standard deviation of the underlying distribution of simple fractional Hamming distance scores. When extreme values are taken from this distribution (the "best match" after several rotations to compensate for head/camera/eye tilt), the mean of this extreme value distribution must obviously be lower as a result of the larger standard deviation[2].

To test this explanation, we rescaled all of the Hamming distances based on the number of bits n used in the computation. Daugman [19] gives a rescaling rule to adjust Hamming distances based on the number of bits used. We used Daugman's SQRT normalization (given in equation 14 of [19]) to rescale all of the Hamming distances. However, we replaced his scaling parameter with our own parameter, 4929, based on the average number of bits used in a comparison between two iris codes for our algorithm and dataset:

$$\mathrm{HD}_{\mathrm{norm}} = 0.5 - (0.5 - \mathrm{HD}_{\mathrm{raw}})\sqrt{\frac{n}{4929}}. \tag{1}$$

The raw Hamming distance means and the SQRT normalized means are shown graphically in Figure 12. Before the SQRT normalization, the difference in mean between set 1 and set 3 was 0.4483 - 0.4265 = 0.0218. After SQRT normalization, the difference was .4475 - .4305 = 0.0170. Therefore, SQRT normalization accounted for nearly a quarter of the difference in the means between set 1 and set 3. It is possible that if our eyelid/eyelash detection software were better, SQRT normalization would account for a larger portion of this difference.

A second possible explanation for the decrease in the mean of the non-match distributions relates to the degrees of freedom in a binomial distribution. The probability mass function for a binomial distribution is given by

$$P(X = x) = \begin{pmatrix} n \\ x \end{pmatrix} p^x (1-p)^{n-x}, \tag{2}$$

where p is the probability of success (of "heads" in a toin-coss, or of a "1" in an iris code bit), and n is the number of independent degrees of freedom (independent Bernoulli trials). Non-match Hamming distance scores follow a binomial distribution with $p = .5$. One would expect a value of n = "number of bits"

[2]We would like to thank one of our anonymous reviewers for suggesting that we explore this idea.

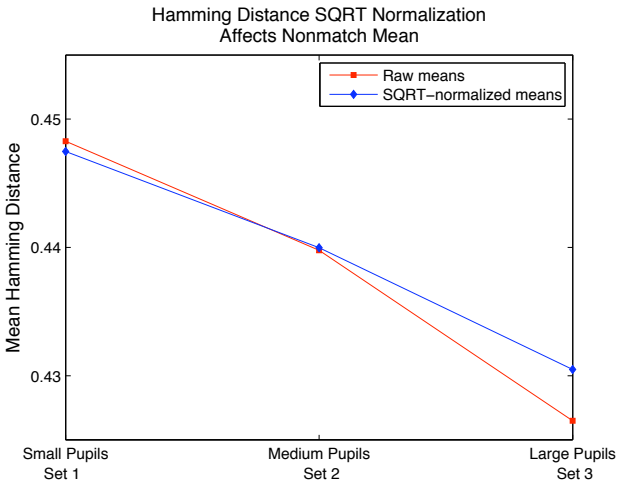

Figure 12: SQRT normalization accounts for part of the decrease in means between set 1 and set 3.

for the second parameter. However, bits in any given iris code have internal correlations arising from iris features [19]. Because of these correlations, the number of independent degrees of freedom is significantly smaller than the actual number of bits in an iris code. In [7], Daugman reports that for his 2048-bit iris codes, his non-match distribution follows a binomial distribution with $n = 249$.

The number of degrees of freedom is significant, because a distribution with more degrees of freedom has a smaller standard deviation. Therefore, a possible explanation for the phenomenon in Figure 11 is that there are fewer degrees of freedom in the underlying distribution for set 3 (large pupils). The simple fact that there is less iris area available in the images may result in fewer degrees of freedom in the distribution.

Up to this point, we have displayed the match and nonmatch distributions separately. In order to summarize both distributions in one figure, we plotted the performance for all three sets in a Decision Error Threshold (DET) curve (Figure 13). This DET curve shows how the false accept rate and false reject rate in our system would change as we vary the decision threshold. Since we aim to minimize error rates, we want the performance curves to come close to the lower left corner of the graph. This figure reiterates what we have already explained in the preceding paragraphs; namely that the performance of the system on the dataset containing small pupils is clearly better than the performance on the datasets with medium or large pupils.

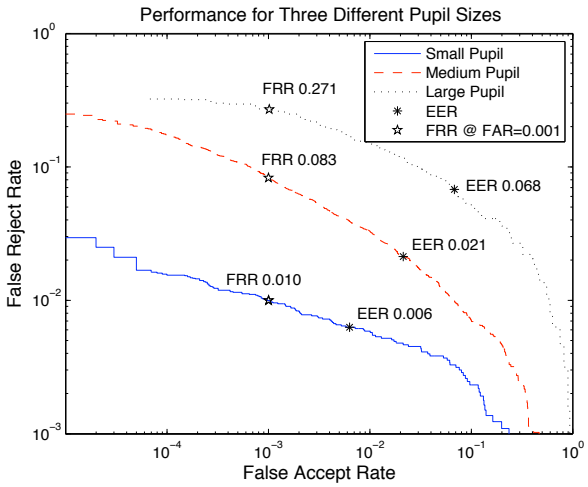

Figure 13: Our iris biometric algorithm performs significantly better when the dataset contains eye images with small pupils.

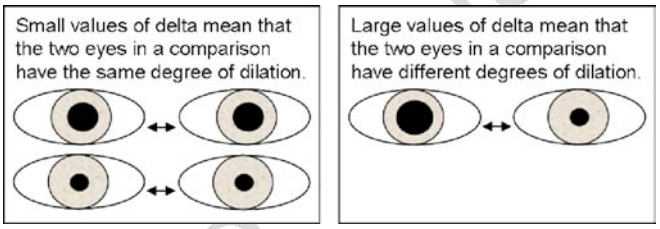

Figure 14: Δ is a quantity that refers to the difference in dilation between two eyes in a comparison.

6 Comparisons between iris images with varying degrees of dilation

Our next experiment deals with comparisons in which the two irises being compared have varying degrees of relative dilation. In order to discuss varying degrees of dilation, we defined a quantity, Δ, which measured the difference in dilation between two eyes (Figure 14):

$$\Delta = \frac{\text{pupil radius}_1}{\text{iris radius}_1} - \frac{\text{pupil radius}_2}{\text{iris radius}_2}. \tag{3}$$

For a comparison in which one eye has a very small pupil, and one eye has a very large pupil, Δ is very large. Within the comparisons done in our experiments, we measured a Δ value as large as 0.487. A comparison between two eyes taken under identical lighting conditions could have $\Delta = 0$.

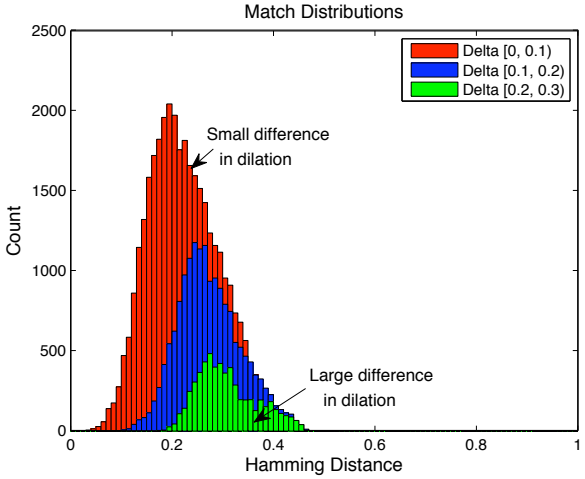

Figure 15: The match distribution of comparisons between eyes with different degrees of dilation has a larger mean than the distribution of comparisons between eyes with the same degrees of dilation.

Again, we separated all comparisons into sets. For this experiment, the first set contained comparisons with small values of Δ. The last set contained comparisons with large Δ, or comparisons between two eyes of very different degrees of dilation. We ran this experiment for both the match distribution (same subject) and the non-match distribution (different subjects). The distributions are shown in Figures 15 and 16. For the match distribution, there were almost no comparisons with Δ in the 0.3 to 0.4 range, so Figure 15 only shows three sets. The non-match distribution had more comparisons in that range, and therefore, four sets are shown.

The match distribution shifts to the right as the difference in pupil dilation increases. The mean of the distributions for comparisons with Δ between 0 and 0.1 is 0.2250. The mean of the distribution for comparisons with Δ between 0.1 and 0.2 is 0.2747. When Δ is between 0.2 and 0.3, the mean is 0.3084. In contrast, the non-match distribution does not shift noticeably. The means of the non-match distributions do not show either an increasing or decreasing trend. The means of the distributions were, in order from smallest Δ to largest Δ, 0.4420, 0.4432, 0.4439, 0.4438.

The Decision Error Threshold (DET) curve for our second experiment is shown in Figure 17. This graph shows that performance rates are higher when we only consider comparisons between eye images of similar dilation ratios.

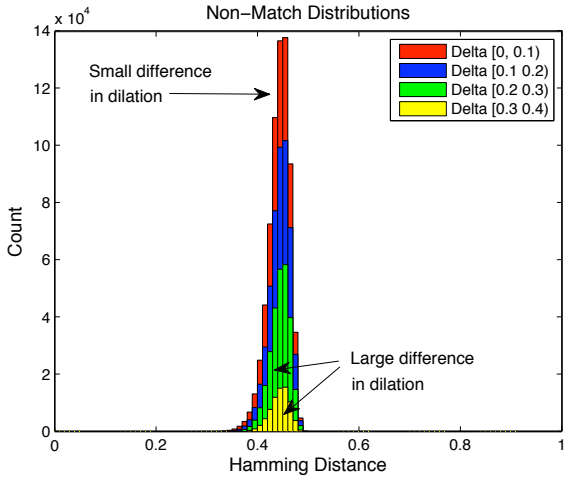

Figure 16: The non-match distribution is not noticeably affected by the difference in dilation in the comparisons.

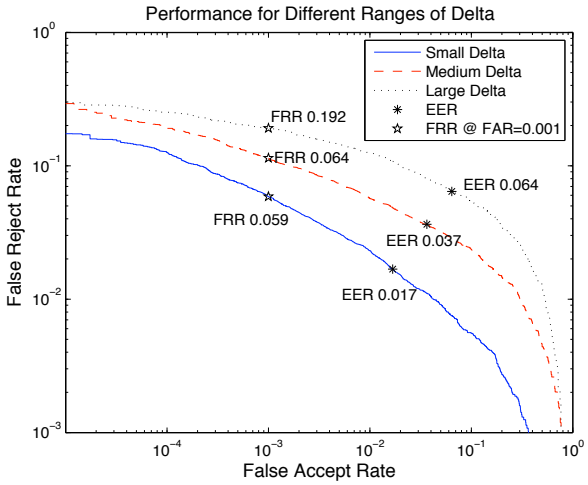

Figure 17: Our iris biometric algorithm performs better when there are not large differences in pupil sizes.

7 Conclusion

We have quantitatively characterized the effect of pupil dilation on iris recognition performance. We found that comparisons between two dilated eyes followed a distribution with a mean fractional Hamming distance of 0.06 higher than the mean of the distribution for non-dilated eyes. The means of both the match and the non-match distributions are expected to fall between 0 and 0.5. Therefore, a shift of 0.06 is nontrivial, amounting to twelve percent of this range. The difference in performance may be partially due to the fact that there is less iris area available in an image of a dilated eye. Points on an eye image are also sampled more closely together for a dilated eye image as compared with a non-dilated image.

We further found that the difference in dilation between an enrollment image and an image to be recognized has a marked affect on the comparison. Comparisons between images with widely different degrees of dilation follow a distribution with a mean about 0.08 higher than the mean of the distribution of images with similar degrees of dilation.

Based on our results, we recommend that a measure of pupil dilation should be created as meta-data to be associated with each generated iris code. This would allow systems to characterize the reliability of an iris code match as a function of the pupil dilations in the underlying images.

One possible line of future work suggested by our results concerns pre-processing the iris image to create an iris code. When the degree of pupil dilation is large, so that the width of the iris is small in pixels, it may be worthwhile to include a super-resolution step in the pre-processing.

8 Acknowledgements

Iris biometrics research at the University of Notre Dame is supported by the National Science Foundation under grant CNS01-30839 and by the Central Intelligence Agency. 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.

The authors also thank Patrick Grother from NIST, and two anonymous reviewers, for their helpful suggestions.

References

[1] Iridian Technologies. http://www.iridiantech.com/, accessed Jan 2008.

[2] LG. http://www.lgiris.com/, accessed Jan 2008.

[3] Sagem Morpho. http://www.morpho.com/, accessed Jan 2008.

[4] Securimetrics. http://www.securimetrics.com/, accessed Jan 2008.

[5] K.W. Bowyer, K.P. Hollingsworth, and P.J. Flynn. Image understanding for iris biometrics: A survey. *Computer Vision and Image Understanding*, 110(2):281–307, 2008.

[6] John Daugman. High confidence visual recognition of persons by a test of statistical independence. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15(11):1148–1161, November 1993.

[7] John Daugman. How iris recognition works. *IEEE Transactions on Circuits and Systems for Video Technology*, 14(1):21–30, 2004.

[8] Harry Wyatt. A minimum wear-and-tear meshwork for the iris. *Vision Research*, 40:2167–2176, 2000.

[9] Li Ma, Tieniu Tan, Yunhong Wang, and Dexin Zhang. Efficient iris recognition by characterizing key local variations. *IEEE Transactions on Image Processing*, 13(6):739–750, June 2004.

[10] Jason Thornton, Marios Savvides, and B.V.K. Vijaya Kumar. A Baysian approach to deformed pattern matching of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 29(4):596–606, April 2007.

[11] Zhuoshi Wei, Tieniu Tan, and Zhenan Sun. Nonlinear iris deformation correction based on Gaussian model. In *Springer LNCS 4642: Int. Conf. on Biometrics*, pages 780–789, Aug 2007.

[12] Xiaomei Liu, Kevin W. Bowyer, and Patrick J. Flynn. Experiments with an improved iris segmentation algorithm. In *Fourth IEEE Workshop on Automatic Identification Technologies*, pages 118–123, October 2005.

[13] Karen Hollingsworth. Sources of error in iris biometrics. Master's thesis, University of Notre Dame, 2008.

[14] Karen Hollingsworth, Kevin Bowyer, and Patrick Flynn. All iris code bits are not created equal. In *Biometrics: Theory, Applications, and Systems*, Sept 2007.

[15] Richard P. Wildes. Iris recognition: An emerging biometric technology. *Proceedings of the IEEE*, 85(9):1348–1363, September 1997.

[16] Libor Masek and Peter Kovesi. MATLAB source code for a biometric identification system based on iris patterns. The University of Western Australia, 2003. http://www.csse.uwa.edu.au/~pk/studentprojects/libor/.

[17] B. Winn, D. Whitaker, D.B. Elliott, and N.J. Phillips. Factors affecting light-adapted pupil size in normal human subjects. *Investigative Ophthalmology and Visual Science*, 35:1132–1137, 1994.

[18] ISO/IEC Standard 19794-6. Information technology - biometric data interchange formats, part 6: Iris image data. Technical report, International Standards Organization, 2005.

[19] John Daugman. New methods in iris recognition. *IEEE Transactions on Systems, Man and Cybernetics - B*, 37(5):1167–1175, Oct 2007.