Detection of Iris Texture Distortions By Analyzing Iris Code Matching Results

Sarah Ring and Kevin W. Bowyer
Department of Computer Science and Engineering, University of Notre Dame

Abstract— **Identity verification using iris biometrics is degraded if portions of the iris texture have a distorted appearance in the image. Distorted texture appearance can result from various causes, including contact lens artifacts, occlusion by hair, eyelids and eyelashes, and specular highlights. Previous work in iris biometrics attempts to cope with occlusion by eyelids / eyelashes and with specular highlights through improved segmentation of the iris region. Our approach assumes that some local distortions of the iris texture are not detected at the segmentation stage, and that these generate corresponding regions of local distortion in the iris code derived from the image. We introduce an approach to detect such regions of local distortion in the iris code through analysis of the iris code matching results. We know of no previous work that attempts to detect distortions of iris texture through analyzing the iris code matching results.**

I. Introduction

Research in iris biometrics has expanded dramatically in recent years [3,12]. One major current theme involves maintaining high accuracy while dealing more flexibly with larger and broader user populations. The research reported here can be seen as fitting within this theme.

Various factors can cause the imaged iris texture to contain local distortions of the true texture. A local distortion of the imaged iris texture causes the bits of the iris code that correspond to the local area to be randomly changed from their true value.

When images of two different irises are matched, local distortions to the iris texture in one or both images are relatively unimportant. This is because the iris codes from two different irises are expected to match only at a random level, and so randomizing the bits in an area of one or both iris codes does not affect the expected result. However, when two images of the same iris are matched, local distortions to the texture in one or both images increases the chances of a false non-match.

This paper presents an approach to automatic detection of local distortions of iris texture in images of the same iris. Existing approaches to detecting local distortions in the iris texture focus on analysis of the iris image. In contrast, our approach to detecting local distortions to the iris texture in images of the same iris focuses on analysis of the iris code matching results. This 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.

Section 2 discusses local distortions of the imaged iris texture and presents some example images. Section 3 outlines how such local distortions in the iris texture affect the iris code, and how this in turn affects the results for matching two images of the iris. Section 4 presents an algorithm for isolating regions of the iris code matching results that correspond to local distortions in one of the underlying iris images. Lastly, Section 5 discusses the results and future work.

II. Local Distortions of Iris Texture

Local distortions of the imaged iris texture can result from a variety of causes. Image artifacts introduced by wearing contact lenses is one class of causes. It is important to distinguish between regular and "cosmetic" contact lenses. It is known that cosmetic lenses may cause drastic distortion of the imaged texture [4]. However, it is less widely appreciated that normal, clear contact lenses can and do cause small, local distortions to the imaged iris texture.

Some examples of artifacts resulting from contact lenses are shown in Figure 1. These images are from the Iris Challenge Evaluation (ICE) dataset [11]. Images in this dataset were acquired using an LG 2200 iris biometric system [9]. One specific type of artifact 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 1(a).This particular logo, the "AV," is functionally important in that it is intended to help contact lens wearers to put the lens on correctly. Contact lenses can also result in a light or a dark circular outline in the image. These outlines represent the boundary between the outer portion of the lens that simply rests on the eye surface and the inner part of the lens that creates the prescribed vision correction. This type of artifact is also visible in the images in Figure 1(a). The intensity of this artifact may be related in some way to the prescription strength of the lens.

Still another type of contact-lens-based artifact can result from an ill-fitting lens. An example of this type of artifact is visible in the images in Figure 1(b) as two concentric white arcs that effectively are specular highlights, and a further concentric black arc that is effectively a cast shadow. The artifact here is apparently caused by a portion of the rim of the contact lens not being properly seated on the surface of the eye. The subject in this case reported that the contact lenses worn at the time did not fit well on their left eye.

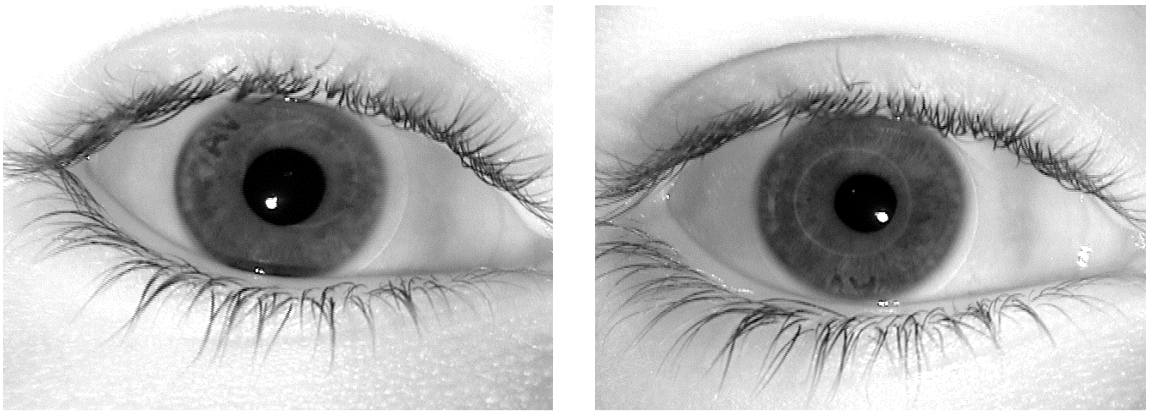

(a) Two images of the same iris (04855d115 and 04855d227), with the “AV” printed on the contact lens overlaying different parts of the iris.

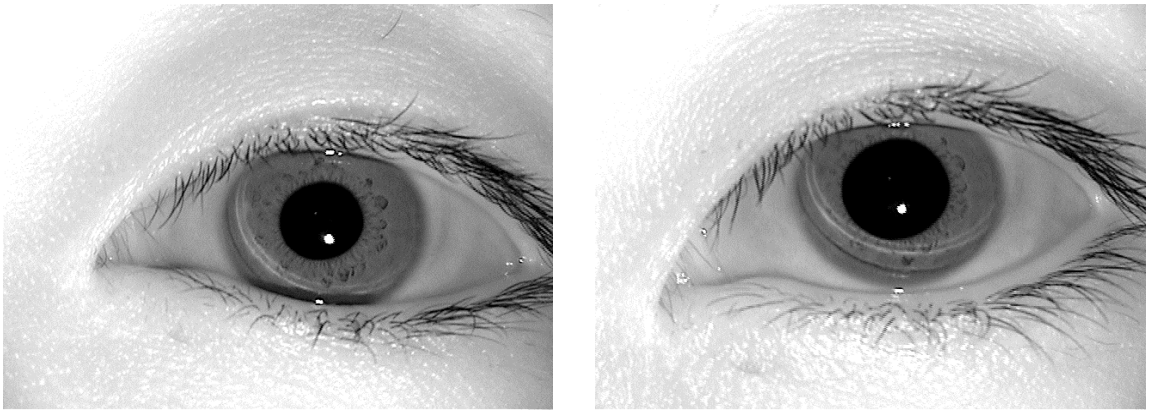

(b) Two images of the same iris (04221d1060 and 04221d122), with artifacts due to poor contact lens fit overlaying different parts of the iris.

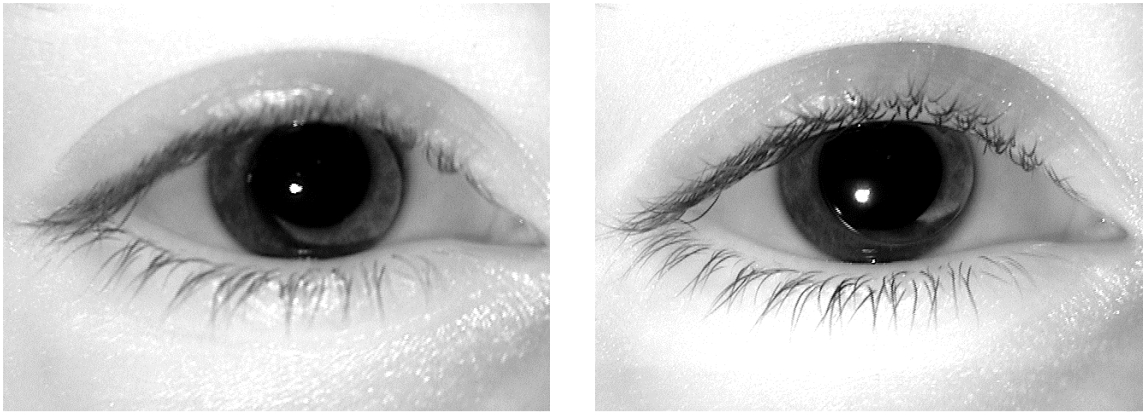

(c) Two images of the same iris (04727d13 and 04727d22), with artifacts due to “hard” contact lens.

Fig. 1. Example iris images in which contact lenses result in significant artifacts that distort the imaged iris texture.
In (a), note the AV in the upper left of the iris on the left and in the lower center of the image on the right; also note the faint white circle that extends to the upper right of the pupil in the left image, and to the left of the pupil in the right image. In (b), note the two concentric white arcs and the dark arc that appears toward the lower left of the pupil in both images. In (c), note the differently-placed distortions on the right side of the iris. Such artifacts can result in local regions of iris texture that do not match, thereby increasing the chances of a false non-match. The two images in (a) are used as a running example in this paper to illustrate our approach.

Other types of contact lens artifacts beyond those illustrated in Figure 1 occur, but examples are not included here due to space limits. We are not aware of any source that tracks the relative frequency of these different types of contact lenses, lens features and lens fitting problems in the general population. However, it seems reasonable to assume that the frequency of strong prescription strength lenses increases with age, and that functionally important features such as the "inside out" logo are likely to become increasingly common.

Note that the contact lens artifacts shown in Figure 1cannot be detected in the way that Daugman has outlined for detecting cosmetic contact lenses [4]. Daugman suggests detecting the presence of cosmetic lenses that are produced using a dot-matrix style printing process through analysis of the power spectrum in the Fourier space. The artifacts depicted in Figure 1 are not the result of a dot-matrix type printing process. Also, the artifacts in Figure 1 are localized in the lens and distort only a portion of the iris texture.

Local distortions of the imaged iris texture can also be caused by inaccurate segmentation. Iris images can present challenging segmentation issues. Among these challenges are regions of the iris that are partially occluded by eyelashes and regions of iris texture that are occluded by inter-reflections. In principle, these occlusions of the iris texture are detected in the segmentation process and marked in the mask that goes with the iris code. However, local occlusions of this type that are not found in the segmentation step can still be detected by our approach.

III. Effects of Local Distortions on Iris Code Matching Results

The typical approach to iris biometrics involves representing the iris texture observed in an iris image by a binary iris code [3,5]. The iris code results from first segmenting the iris region of the image, then creating a normalized iris image, then applying a texture filter at each of a set of sample points in the iris image, and finally quantizing each of the real and imaginary parts of the texture filter result to a 0 / 1 value based on whether it is negative / non-negative.

Figure 2(a) shows the iris segmentation results for the images in Figure 1(a) as obtained using Liu's [10] refinement of the ICE open source iris biometrics software [11]. The white circles represent the pupil-iris and the iris-sclera

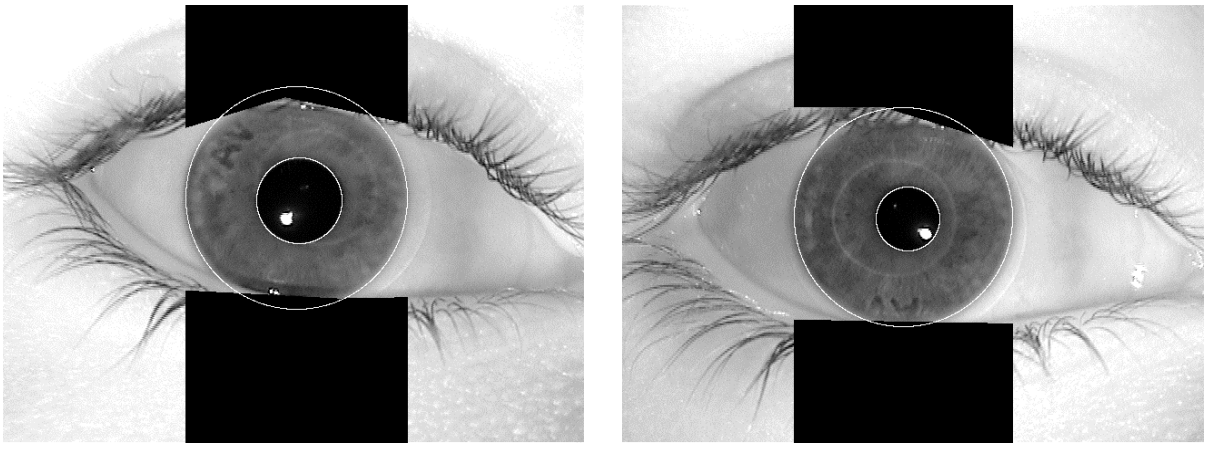

(a) Illustration of the segmentation results for the two iris images in Figure 1(a).
In each image, the inner white circle shows the detected pupil-iris boundary, the outer white circle shows the iris-sclera boundary, and the black regions show the detected eyelid occlusion boundary.

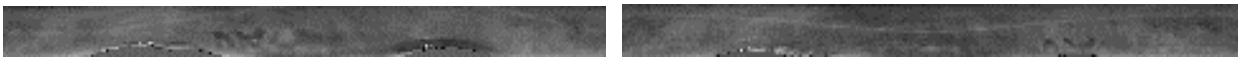

(b) The normalized iris images corresponding to the iris regions found in (a).

(c) The segmentation masks corresponding to the normalized images in (b).

Fig. 2. Illustration of iris image segmentation results, normalized images, and segmentations masks for the images in Figure 1(a).
In (a), note the AV in the upper left of the iris on the left and in the lower center of the image on the right. Also in (a), note the faint white circle that extends to the upper right of the pupil in the left image, and to the left of the pupil in the right image. In (b), note the AV and the white circle are still visible in each normalized image. The white circle in (a) appears as a curve that runs from left to right in the image in (b). In the two images in (b), the white curves wind through the same general area of the upper right of the images. Part (c) depicts the segmentation masks, where black indicates the region of the iris not occluded by the eyelids and white indicates occluded regions. Note that the segmentations are plausible but not perfect. The segmentation does not quite get all of the upper eyelid in either image. In the image on the left, while the segmentation does get nearly all of the lower eyelid, there is also a shadow cast on the iris that can distort the texture filter results and that is not found by the segmentation.

boundaries found by the system. The black blocks at the top and bottom of the iris represent the eyelid-iris occlusion boundaries found by the system.

Based on the iris segmentation, a normalized image and a segmentation mask are created. The normalized image is often informally referred to as an "unwrapped" or a "rectangular" iris image. Figure 2(b) shows the normalized image and 2(c) shows the corresponding segmentation mask for the two example images. The segmentation mask indicates the portions of the normalized iris image that are not considered to generate valid bits of the iris code because those portions are occluded by eyelid / eyelash.

A texture filter is applied to the normalized iris image at a pre-defined set of points. In the system used here, the normalized image is 20x240, and a 1D log-Gabor filter is applied at each point. The points at which the filter is applied can be viewed as sampling at increments along the radial distance between the pupil-iris boundary and the iris-sclera boundary and at increments of angular distance around the iris. At each point that the filter is applied, a complex-valued result is obtained. The real part and the imaginary part of each result are each quantized to 0/1, giving two bits of iris code for each texture filter result. Thus the system used here generates an iris code that is 2 x 20 x 240 = 9600 bits.

At the stage that the iris code is generated, the mask should be augmented to take into account iris code bits that are likely to be inconsistent, or fragile [2,6,7]. Essentially, if the real or the imaginary part of the texture filter result is close to 0, then random variation between images of the same iris could easily cause that value to flip between negative and non-negative, or between 0 and 1 in the iris code. For this reason, the iris code bits that correspond to the smallest-magnitude filter values are added to the mask. The mask then represents a combination of the regions of the iris that the segmentation stage found to be occluded, and the bits of the iris code that the texture encoding stage determines to be likely to be inconsistent. For the results presented here, one quarter of the iris code bits corresponding to the smallest-magnitude filter values are masked. Thus, even if the iris appears in the image with no occlusion, the number of iris code bits actually used is 4,800 – 1,200 = 3,600 corresponding to the real parts of the texture filter results and another 3,600 corresponding to the imaginary parts. Occlusion of the iris results in a larger number of bits masked out and so a smaller number used for matching.

In general, matching of two iris codes is done by circularly shifting one relative to the other over some range, computing the normalized Hamming distance for each shift, and taking the shift that results in the best match. The point of the shifting is to account for possible relative rotation of the iris in the two images. This step is essential for systems that image a single eye at a time. Systems that image both eyes at once may be able to avoid or minimize this step.

The normalized Hamming distance is the fraction of the bit positions at which the two iris codes do not agree, and it is computed only over those bit positions that are unmasked for both iris codes. The normalized Hamming distance ranges between 0, indicating that the two codes agree at all unmasked bit positions, and 1, indicating that the two iris codes do not match at any of the unmasked bit positions. The expected normalized Hamming distance between images of two different irises is 0.5. The expected distance between two images of the same iris is significantly smaller than 0.5 but is not zero. The actual expected value for two images of the same iris depends to a large degree on how well the image acquisition conditions are controlled. The probability of making identity verification errors when matching iris codes can be estimated by analyzing the match and non-match distributions obtained from a body of experimental data; e.g., see [5].

Image artifacts and segmentation inaccuracies have a generally asymmetric effect on the non-match and match distributions. For a non-match comparison, theory predicts that the a priori probability of a match at any given bit position is 0.5, and experimental evidence has broadly confirmed this. (Taking the best circular shift to account for possible relative rotation does reduce the expected value slightly.) Thus, a distortion of the iris texture or a segmentation inaccuracy in one or both images, which can be assumed to randomize the corresponding iris code bits, is not expected to affect the mean or variance of the non-match distribution. For a match comparison, a distortion of the iris texture or a segmentation inaccuracy randomizes iris code bits that otherwise have a high probability of matching. Thus, such problems tend to increase the mean of the match distribution and to give it a "tail" that trends toward the non-match distribution. Thus the practical effect of texture distortions or segmentation inaccuracies is that, for a given threshold for distinguishing between non-match and match, the chances of a false non-match result are increased. Alternatively, if the chances of a false non-match are to be held constant, the threshold must be changed so that the chances of a false match are increased.

IV. Algorithm for Detecting Local Distortions of the Iris Texture

We wish to be able to detect small, localized distortions of iris texture or segmentation inaccuracies that arise from any of a variety of causes. One important insight is that there is no need to detect such problems in the case of a non-match comparison. A second important insight is that if such a problem occurs in one or both of the images in a match comparison, it generates a simple and recognizable feature in the iris code match results. Consider the iris code match results as an image, where gray indicates a position that was masked out in one or both of the images being matched, white indicates matching bits, and black indicates non-matching bits. A local region in the iris code match results that has a high density of non-match 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 image with a large number of small windows.

2. Compute the normalized Hamming distance separately for each window.
3. Identify as "outlier" windows those that have an unusually high concentration of non-match bits.

The normalized Hamming distance for the match can then be re-computed without the contribution from the outlier windows. This has the effect of "correcting" the tail to the high side of the match distribution and reducing the chances of a false non-match.

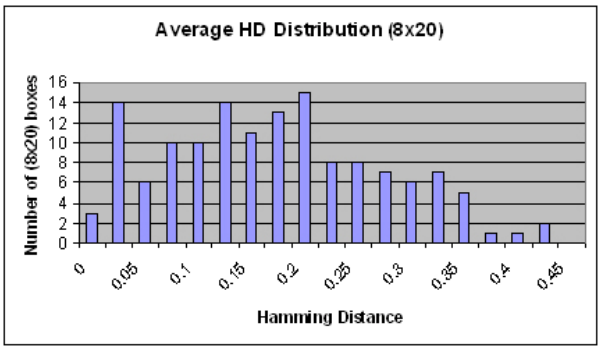

Fig. 3. Distribution of Normalized HD Across Windows.
The results image is covered with 8x20 windows that overlap by one half in each dimension, and the normalized Hamming distance is computed for each window.

To explore this idea further, we consider the match results for the iris images shown in Figure 2. The matching results can be considered as a pair of 20x240 images, one for match bits that derive from the real value of the texture filter and one for match bits that derive from the imaginary value. We cover each 20x240 image with windows of size 8x20 that overlap by one-half in each dimension. Thus there are 4x23 = 92 windows for each of the two images. Windows that have too few unmasked bit positions (in this example, less than 20) are dropped from consideration, as the statistics from a small number of match bits are too noisy. Figure 3 shows the resulting distribution of normalized Hamming distance across the windows. The normalized Hamming distance for the overall match is 0.21. This distribution of normalized Hamming distance for the 8x20 windows ranges from a low of nearly 0 to a high of nearly 0.5.

We explore this further by examining the locations of the windows that have the highest normalized Hamming distance. Figure 4 (b) shows the locations of the ten windows with the highest normalized Hamming distance, for each of the real and the imaginary parts of the results. Here, gray indicates a masked bit position, white indicates a matching bit position, and black indicates a non-matching bit position. Figure 5(c) shows the location of the twenty windows drawn on the two iris images. The location of the AV logo in each iris image is covered, as is the location of the shadow near the lower eyelid in the image on the left in Figure 2(a)-(b). The window in the upper right of the iris images in Figure 5(c) contains an area that has the curved white line in each image running through it. There are no windows in the area of the upper eyelid occlusion because the union of the occlusion masks for the two images left only a very small region of

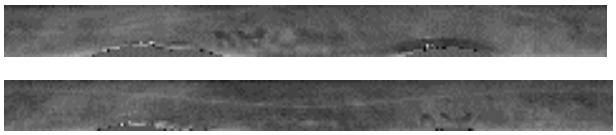

(a) The normalized iris images from Figure 2.

(b) The ten windows of the real (above) and of the imaginary (below) match results that have the highest normalized Hamming distance.

(c) Locations of the twenty windows identified in (b) drawn on each of the iris images shown in (a).

Fig. 4. Regions of Texture Distortion / Segmentation Inaccuracy.
Note that the windows cover the region of the printed AV logo in each image, the region of the shadow on the iris on the left in Figure 2(a)-(b), and a region in the upper right in which the curved white line in each image runs through the same window.

segmentation inaccuracy. The normalized Hamming distance for the match with these windows removed from the calculation drops from 0.21 to 0.15.

V. Discussion and Future Work

There are various types of artifacts that can occur in iris images that increase the chances of a false non-match. We have shown several examples, including artifacts associated with wearing contact lenses and shadows. Artifacts arising from specular reflections and segmentation inaccuracies can cause similar problems. Because these artifacts have a variety of different and subtle appearances in an iris image, the task of detecting them through analysis of the iris image is extremely difficult.

An artifact in an iris image or an error in the iris segmentation increases the chances of a false non-match because it randomizes the corresponding bits of the iris code. When this iris code is then matched against another iris code for the same iris, the effect is a local area of randomized match results. While the match as a whole might have a normalized Hamming distance of 0.2 or lower, the local region in the match results corresponding to the artifact in the iris image might easily have a normalized Hamming distance of nearly 0.5, as shown in the example here.

We have outlined a method of detecting the effect of an image artifact in the iris code matching results, rather than directly detecting the artifact in the iris image. The problem is simpler when attacked this way because the various types

of artifacts that can occur in the iris image cause the same type of artifact in the matching results. In the matching results for two images of the same iris, a local area of near-random match results indicates an artifact in one of the matched images. Once local areas of outlier match results are identified, the normalized Hamming distance for the match can be corrected by re-computing without the problem areas.

In practice, this approach would need to be applied only in cases where the initial normalized Hamming distance indicates a borderline match:

1. Compute normalized Hamming distance from the two iris codes, $HD_{initial}$.
2. If $HD_{initial} > T_{high}$, then declare non-match.
3. If $HD_{initial} < T_{low}$, then declare match.
4. If $T_{low} <= HD_{initial} <= T_{high}$
    a. Cover results image with small windows.
    b. Compute normalized HD for each window.
    c. Identify outlier windows.
    d. Re-compute normalized HD with outlier windows removed, $HD_{corrected}$.
5. If $HD_{corrected} > T_{final}$, declare non-match, else declare match.

Parameters of this approach include the size of the windows, the pattern of covering the results image with windows, and the minimum number of unmasked result bits needed to compute a result for the window. The size of a window and the threshold on the number of unmasked bits needed should be large enough to compute a statistically reliable estimate of the normalized Hamming distance. However, it should not be too much larger than the size of the image artifacts, so that the normalized Hamming distance for the artifact area is not diluted by correct match results. The pattern of covering the results image with windows should be dense enough that any significant artifact is likely to be mostly covered by some window.

The approach introduced here is likely to become increasingly important as iris biometrics is applied to large-scale populations in real-world applications. It is also likely to become increasingly important as the variety of functional features in contact lenses increases. There is even current research on embedding circuits and displays in contact lenses [13].

There are several important topics of future work to be considered. One is to verify that this approach improves the match distribution on a larger experimental dataset of images such as those in Figure 1. Another possible topic of future work is to determine the extent to which this approach might be used as a primary method of detecting eyelash and eyelid occlusion.

An important practical element of future work is to develop an automatic, adaptive method of determining the region of match results that has been corrupted by texture distortions. This would remove the need to set parameters for the size and number of windows covering the results image. The window size used for the experimental results in this paper was motivated by the approximate size of the "AV" logo as it appears in sample images.

A more speculative line of research involves whether or not this approach might help in the matching of iris images in which there a large difference in pupil dilation [8]. If the texture distortion caused by pupil dilation is not uniform, but has local areas of high distortion, this approach may be of some use.

As iris biometrics moves toward possible use in large-scale applications, it is important that no element of society be differentially disadvantaged in using the technology. It would be especially unfortunate if difficulty in using iris biometrics technology were correlated with a covariate such as wearing contact lenses. The results of our work have the potential to reduce such problems.

ACKNOWLEDGMENTS

Thanks to Patrick Flynn, Karen Hollingsworth, and John Daugman for various helpful discussions during this work,and to the anonymous reviewers for suggestions that improved the content and presentation of this paper.

This work sponsored in part by the Central intelligence Agency and by the National Science Foundation CNS 01-30839. Sarah Ring is an Ateyeh Undergraduate Research Scholar at the University of Notre Dame.

REFERENCES

[1] ACUVUE 2 brand contact lenses. http://www.jnjvision.com/acuvue2.htm

[2] Ruud M. Bolle, SharathPankanti, Jonathan H. Connell, and NaliniRatha. Iris individuality: A partial iris model. *Int'l Conference on Pattern Recognition (ICPR)*, pages II: 927-930, 2004.

[3] Kevin W. Bowyer, Karen Hollingsworth and Patrick J. Flynn, Image understanding for iris biometrics, *Computer Vision and Image Understanding* 110 (2), May 2008, 281-307.

[4] John Daugman, Demodulation by complex-valued wavelets for stochastic pattern recognition, *Int'l Journal of Wavelets, Multiresolution and Information Processing* 1 (1), March 2003, 1-17.

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

[6] John Daugman, personal communication, September 27, 2007.

[7] Karen Hollingsworth, Kevin W. Bowyer and Patrick J. Flynn, The best bits in an iris code, *IEEE Transactions on Pattern Analysis and Machine Intelligence,* to appear.

[8] Karen Hollingsworth, Kevin Bowyer and Patrick J. Flynn, The importance of small pupils: a study of how pupil dilation affects iris biometrics, *Biometrics: Theory, Applications and Systems (BTAS '08)*, September 2008, to appear.

[9] LG Advanced Identity Authentication.`http://www.lgiris.com`.

[10] Xiaomei Liu, Optimizations in iris recognition, PhD Dissertation, University of Notre Dame, 2006.

[11] P. J. Phillips, W. T. Scruggs, A. J. O'Toole, P. J. Flynn, K.W. Bowyer, C. L. Schott, and M. Sharpe, FRVT 2006 and ICE 2006 Large-Scale Results, National Institute of Standards and Technology, NISTIR 7408, http://face.nist.gov, 2007.

[12] Nalini K. Ratha and VenuGovindaraju, editors, *Advances In Biometrics: Sensors, Algorithms and Systems*, Springer-Verlag London Limited, 2008.

[13] Scientists working on embedding circuits, displays in contact lenses, *Computer* 41 (3), March 2008, 22.