Degradation of Iris Recognition Performance Due to Non-Cosmetic Prescription Contact Lenses

Sarah E. Baker, Amanda Hentz, Kevin W. Bowyer, and Patrick J. Flynn

sbaker3,kwb,flynn@cse.nd.edu; ahentz1@gmail.com

Abstract

Many iris recognition systems operate under the assumption that non-cosmetic contact lenses have no or minimal effect on iris biometrics performance and convenience. In this paper we show results of a study of 12,003 images from 87 contact-lens-wearing subjects and 9,697 images from 124 non-contact-lens wearing subjects. We visually classified the contact lens images into four categories according to the type of lens effects observed in the image. Our results show different degradations in performance for different types of contact lenses. Lenses that produce larger artifacts on the iris yield more degraded performance. This is the first study to document degraded iris biometrics performance with non-cosmetic contact lenses.

Index Terms

biometrics, iris recognition, contact lenses, match distribution stability

I. INTRODUCTION

Approximately 28 to 38 million people wear contact lenses in the United States and about 125 million wear them worldwide [14]. Contact lens manufacturing is a growing, multi-billion dollar industry. As iris recognition becomes more widespread, systems must be flexible in dealing with large and varying populations while also maintaining appropriate security levels. The large number of contact lens wearers in the general population requires recognition systems to operate well for individuals wearing varying types of lenses. There has been a widely accepted belief in iris biometrics research that prescription contact lenses have little or no effect on recognition. This belief has been stated as: "Subjects can generally be recognized through eyeglasses or contact lenses. Colored contact lenses do not affect the enrollment/recognition process" [6], "The use of glasses or contact lenses (colored or clear) has little effect on the representation of the iris and hence does not interfere with the recognition technology" [5] and "Successful identification can be made through eyeglasses and contact lenses" [7]. We present empirical results that challenge this belief.

Some prescription contact lenses result in significant artifacts visible in the iris image. These distortions of the iris texture increase the likelihood of two images of the same iris being falsely declared to be a non-match. On visual inspection, other prescription lenses appear to cause no distortion of the iris texture. We explore whether these contacts nonetheless affect recognition results.

Section II discusses related work. Section III outlines our experimental materials and methods as well as our categorization of images based on the type of artifacts apparent in the image. Section IV presents experimental results using the IrisBEE system with active contours, section V presents analogous results using the VeriEye system [13], and section VI presents results using the an implementation provided by Carnegie Mellon University. Section VII presents a limited set of results for cosmetic contact lenses. Section VIII further illustrates the regions of the iris affected by contact lenses. Lastly, section IX discusses these results and future work.

II. RELATED WORK

Cosmetic or patterned contact lenses are known to affect iris recognition performance. It has been shown that the use of such lenses can allow someone on a watch list to evade detection [12]. Daugman presents a method to detect cosmetic contact lenses manufactured with a "dot-matrix"

type process [12]. In his method, Daugman finds peaks in the Fourier power spectrum that correspond to the periodicity of the dot-matrix pattern.

Wei et al. suggests three different techniques to detect "counterfeit irises" (cosmetic lenses) [4]. They propose detecting iris edge sharpness because cosmetic lenses usually have a sharper iris boundary edge. Another method they test uses iris texture classification and iris-texton features. Their third method uses co-occurrence matrices that characterize relationships between neighboring pixels. They show that these methods are effective although somewhat dependent on the data set. The two self-collected data sets that they use are apparently not available to other researchers.

Many iris recognition algorithms account for "hard" or "gas permeable" contacts [25]. Daugman's algorithm detects the boundaries of such contacts and ignores these regions as artifacts [1], [17].

Soft contacts are usually treated as having no effect on verification. However, some Acuvue [16] lenses have an "AV" or "1-2-3" indicator printed on the lenses in order to ensure proper application and insertion. Our research group has previously reported degradation of the Hamming distance and increased false reject rate due to artifacts such as an "AV" printed on contact lenses [19].

To our knowledge, no work outside of our group has been done to investigate the effect of contact lenses without registration indicators or cosmetic patterns. In previous work [23], we showed that regular prescription contact lenses degraded performance significantly. In this paper we have expanded the data set used in [23] in size, in variety of image quality, in segmentation quality and in complexity of contact lenses. Additionally, we present results using a modified IrisBEE system that uses adaptive segmentation [3], [9], [27], the commercial VeriEye system [13], and the Carnegie Mellon's implementation [28].

III. EXPERIMENTAL MATERIALS AND METHODS

This section describes the iris image dataset, the iris biometrics software used, and the categorization of the types of contact lenses and their resulting artifacts on iris images used in this study.

A. Image Data set

All images were acquired using an LG 2200 iris imaging system [2]. Image acquisitions were performed in the same studio, with the same ambient indoor lighting, and using the same protocol as the images in the Iris Challenge Evaluations (ICE) [9], [10], [24]. We visually inspected all images considered for use in this study and rejected any of noticeably low quality. We also rejected images with IrisBEE segmentation results failing to mask eyelid or eyelash occlusion well. While this selection of well-segmented images was subjective, we tried to maintain the same acceptance criteria for all images. We also rejected images that failed to match well due to large rotation of the eye between images. Both the IrisBEE and the VeriEye systems allow a rotation within 15 degrees, so images rotated more than this were rejected as they were considered outside of the normal operating range.

The rotation of an iris within an image was determined by matching the image against all other images of the same iris. The IrisBEE software reports the shift of the bits required for the two images to be in the closest alignment. The images that required more than a 15 degree shift for the majority of the matches were considered to have a rotation of 15 degrees.

We consider 92 subjects who never wore contact lenses at acquisition, 52 subjects who wore the same type of contact lenses for all acquisitions, 32 subjects who wore contact lenses for some acquisitions and did not wear them for others, and 3 subjects who changed the type of contacts they wore between acquisitions. Figure 1 shows example images of a subject who changed the type of contact lens they wore between acquisitions. In all, we consider images from 124 subjects wearing no contact lenses and images from 87 subjects wearing contact lenses during image acquisition, with some subjects appearing in both categories. Our subjects range in age from 19 to 58 with 75% of the subjects between 19 and 25. The gender breakdown is 85 female and 86 male subjects; and the ethnic breakdown is 36 Asians, 6 Hispanics, 122 Caucasian, and 7 subjects not reporting.

For the 124 non contact lens subjects (248 irises), we have 9,697 images with an average of 80 images per subject. For the 87 contact lens subjects (176 irises), we have 12,003 images with an average of 135 images per subject. We visually inspected all images to ensure all those in the contact lens category showed a visible contact lens and all those in the non contact lens category were indeed without any visible lens.
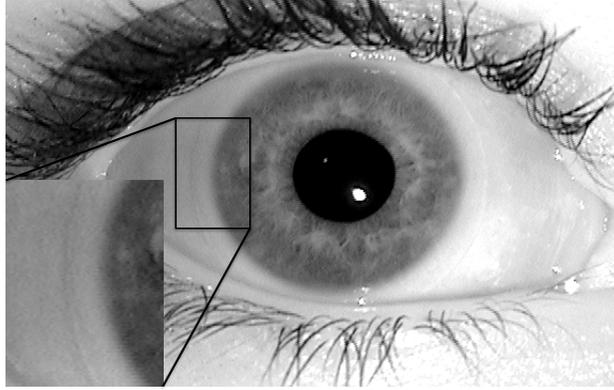
(a) Example of an image for subject 04213 with a contact lens without a circular boundary in the iris. Image name: 04213d445.
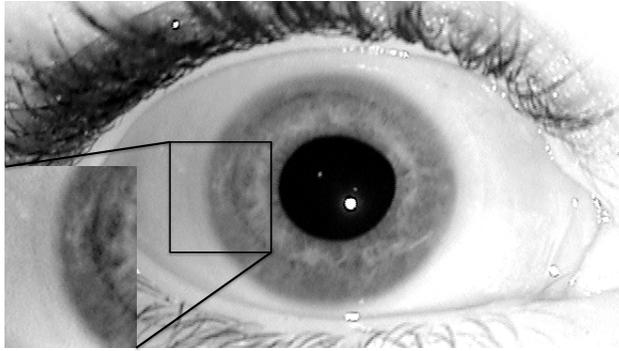
(b) Example of an image for subject 04213 with a contact lens with a circular boundary in the iris. Image name: 04213d451.

Fig. 1.   Example images of a subject who changed the type of contact lens worn in acquisitions.

We used three different iris matching systems to generate match scores for the image data set. First, we used a modified version of the open source IrisBEE system to locate and segment the iris and report match scores as a fractional Hamming distance between 0 and 1 [3], [9]. This version uses a Hough transform to find the pupil and iris regions but uses a balloon-based active contour model to locate the limbic boundaries [27]. Bits in the iris code are masked due to eyelid and eyelash occlusion. Additionally, we mask 25% of the bits as "fragile" or inconsistent [11]. This software accounts for iris rotation of a +/- 15 degree angle by shifting the probe iris code and returning the best Hamming distance found in this range.

Second, we used the commercial VeriEye Iris SDK from Neurotechnology [13]. This system produces match scores on a different scale and with a different polarity than IrisBEE. The

parameter for maximal rotation was set to allow an approximate +/- 14-15 degree angular rotation of the iris.

Third, we used an iris biometric implementation provided by Carnegie Mellon University [28]. This system produces match scores as Hamming distances with the same polarity as IrisBEE.

B. Categorization of Contact Lenses

Our earlier, smaller study [23] notes that the types of contacts used affects match quality differently. We categorized the contact lens images into four categories based upon our visual inspection.

The first category contains those images with a visible contact lens that left no visible artifact on the iris itself (see Figure 2). The boundary of these lenses can often be seen in the sclera, but no artifact is apparent in the iris texture. These lenses are regular soft contact lenses with no indicators to aid in insertion and no correction for astigmatism. We placed 5,867 images from 47 subjects into Category 1.

We placed 3,602 images from 26 subjects into Category 2, consisting of images with a visible contact lens that resulted in a light or dark circular outline within the iris itself as well as an outer outline on the sclera (see Figure 3). Some individuals requiring vision correction for near-sightedness or far-sightedness are also astigmatic (they have corneas that are not perfectly round.) Toric contact lenses have been designed specifically for such individuals [20]. These lenses have two different curvatures, allowing for correction for both astigmatism and "near-" or "far-sightedness." For subjects wearing toric lenses, it is the curvature boundary on these contact lenses that yields the circular outline visible in the iris. Toric lenses have a specific orientation in which they must remain to give proper vision correction. These lenses are often designed to be heavier at the bottom so that the lens will automatically rotate into correct position after blinking or other movement. Because they must maintain the same orientation, toric lenses must be inserted properly, thus they often are marked with indicators on the sides or bottom to aid in insertion. Figure 3(c) shows an example of a toric contact lens. We confirmed with the subject shown in this picture that they wore a toric contact lens in their left eye for every acquisition.

Category 3 consists of images of irises with contacts causing a large visible artifact on the iris. This large artifact may be an "AV" logo (see Figure 4(a)) or numbers, such as "1-2-3" (see Figure 4(c)), printed on the iris for ease of insertion. We have 11 subjects in our data set with
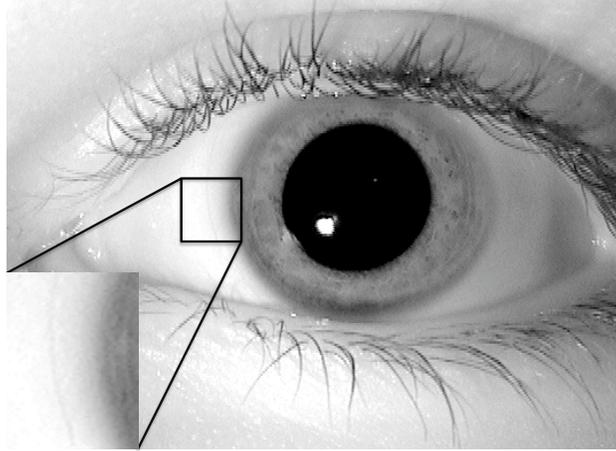
(a) Image name: 04312d593.
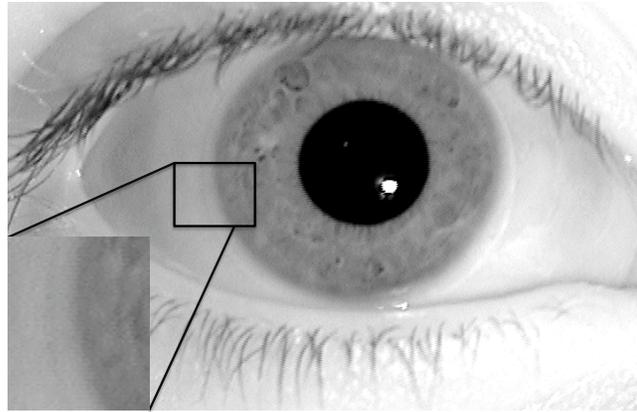
(b) Image name: 04408d479.

Fig. 2. Examples of Category 1 Contact Lens Images. We observe little or no contact lens artifact in the iris region. The arrow points to visible evidence of the contact lens, a faint edge (almost impossible to see) at the border of the contact resting on the sclera.

such markings printed on the lenses. We also placed in this category images with other large artifacts present in their iris such as those in Figs. 4(e) and 4(g). One example of such an artifact results from an ill-fitting lens where a portion of the rim of the lens does not properly rest on the surface of the eye. We ascertained from the subject shown in Figure 4(e) that the contact lenses worn at the time of this image acquisition did not fit well. We classified 1,802 images from 14 subjects as Category three contact lenses.

Finally, we considered the four subjects and 732 images in our data set with rigid gas permeable
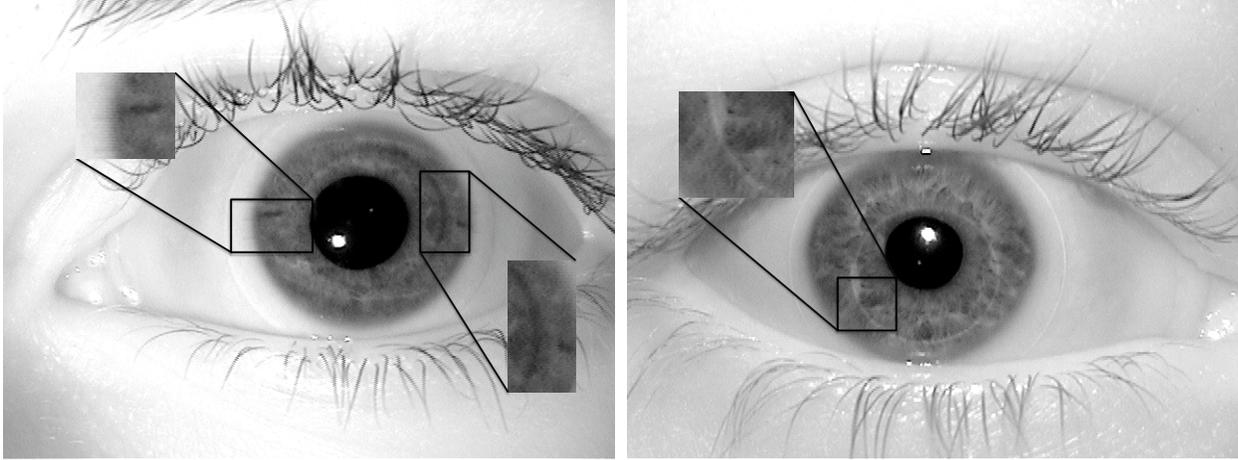
(a) Image name: 04453d767.

(b) Image name: 04613d893.
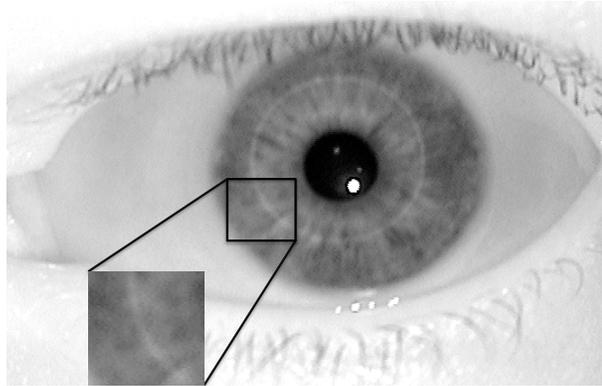
(c) The image is an example image from subject 04261 who confirmed wearing a toric contact lens in their left eye for every acquisition. Image name: 04261d754.

Fig. 3. Examples of Category 2 Contact Lens Images. We observe a circular boundary line of the contact lens in the iris region in addition to the circular boundary on the sclera. In Figure 3(a), we also observe an indicator used for insertion.

[25] contacts (see Figure 5). These contacts are similar to the older "hard" contact lenses. Gas permeable lenses are much smaller in diameter than soft contact lenses, so the entire lens fits within the boundaries of the iris. It is widely accepted that rigid lenses cause a degradation in iris match quality and increase the false reject rate. These images are labeled as Category 4.

Additionally, we placed the 9,697 images from 124 subjects without any contact lenses in Category 0.
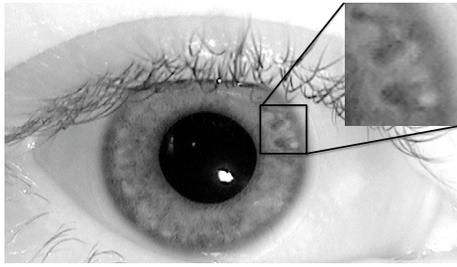
(a) Example iris image with AV on the contact lens. Image name: 04341d696.
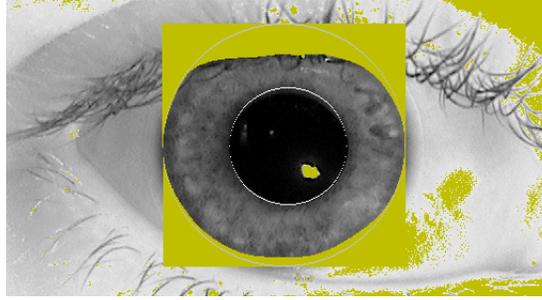
(b) Segmentation of Image 04341d696.
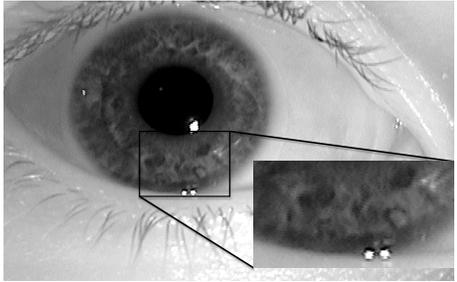
(c) Example iris image with numbers on the contact lens. Image name: 04869d115.
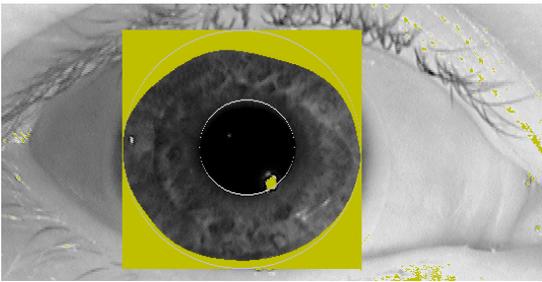
(d) Segmentation of Image 04869d115.
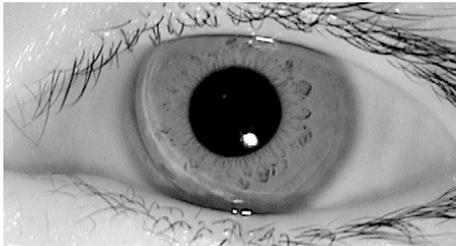
(e) Example iris image with a significant artifact on the iris from the lens. This subject reported this contact lens did not fit his/her eye properly. Image name: 04221d1382.
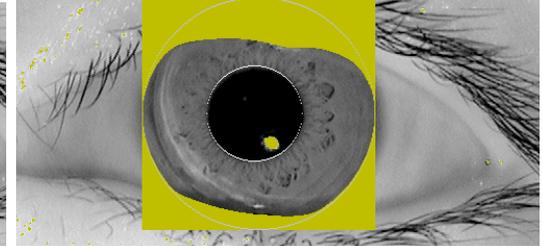
(f) Segmentation of Image 04221d1382.
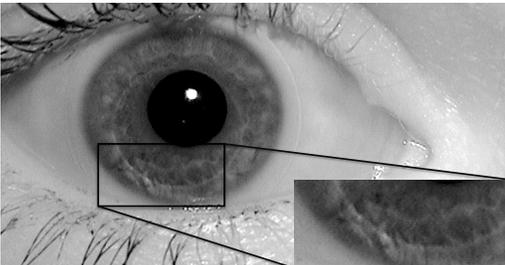
(g) Example iris image with a significant artifact on the bottom left region of the iris from the lens. Image name: 04885d111.
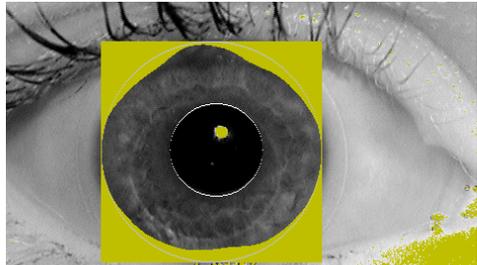
(h) Segmentation of Image 04885d111.

Fig. 4.   Examples of Category 3 Contact Lens Images with the corresponding IrisBEE segmentations.

(a) Image name: 04593d633.

(b) Segmentation of Image 04593d633.

(c) Image name: 04622d629.

(d) Segmentation of Image 04622d629.
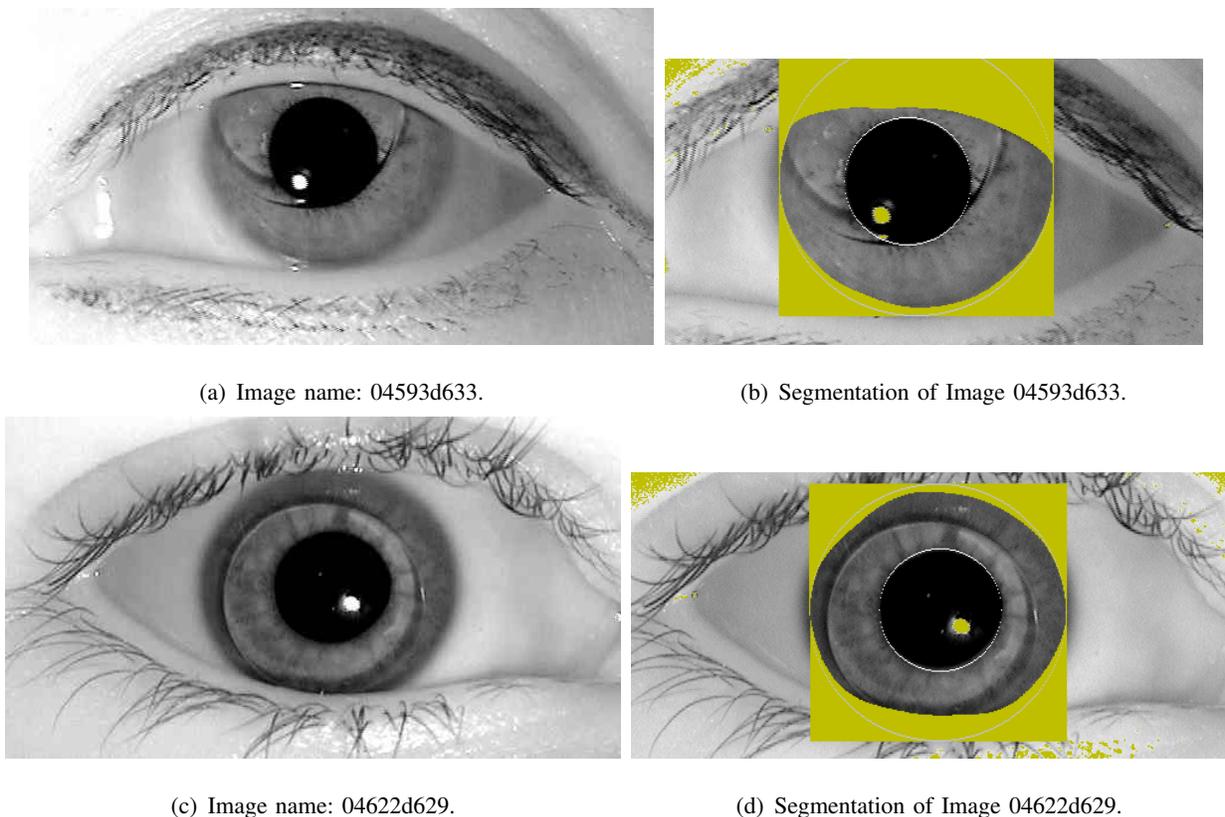
Fig. 5. Examples of Category 4 iris images with hard contact Lenses and the corresponding IrisBEE Segmentations.

C. Experimental Method

We performed an "all versus all" experiment, matching all pairs of images in our data set. We then analyzed the results based on the type of contact lens worn in each image in the pair. We completed this analysis for the IrisBEE with adaptive segmentation, VeriEye, and CMU systems.

D. Cosmetic Lens Study

In addition to our study of prescription lenses, we present results from experiments involving images of subjects wearing cosmetic lenses. We have one cosmetic-contact-lens wearing subject (04780) who regularly attended acquisitions. This subject wore these cosmetic lenses in some sessions, but also wore regular contact lenses in other sessions (see Figure 6). The cosmetic lenses worn by subject 04780 are a pad-printed [22] type of cosmetic contact lenses. This type of lens is a colored contact lens with an uncolored pupillary region and a multicolored pattern designed to cover at least 80% of the subject's iris. Manufacturers of these lenses use several
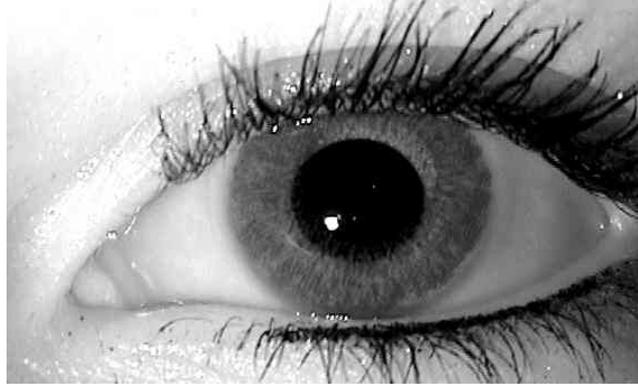
(a) Subject with cosmetic contact lens. Image name: 04780d133.
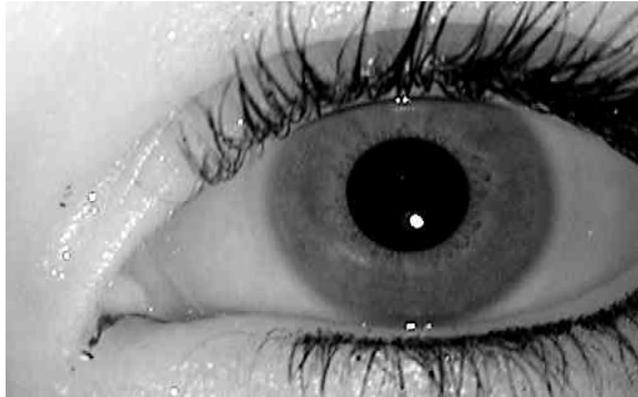
(b) Subject without cosmetic contact lens. Image name: 04780d140.

Fig. 6.   Example of subject 04780 with and without a cosmetic lens

different layers of color in order to simulate the texture of a genuine iris. For example these lenses are designed with a collarette band surrounding the pupil area with jagged edges in order to resemble a true collarette. To further improve the realistic duplication of a human iris, some of the region of the lens covering the iris is left uncolored, which allows for light exchange and interaction between the uncolored openings and the iris itself. Contact lens developers have found that leaving approximately 20% of the iris region uncolored allows for the true iris texture to be visible and helps to make the lens appear more realistic [22].

We also acquired images from two different subjects wearing cosmetic lenses with a dot-matrix pattern as well as images of these subjects not wearing the lenses. The dot-matrix pattern appears as a matrix of colored opaque pigmented dots on the lens, but unlike the pad-printing

cosmetic lenses, the lens is not designed specifically to recreate a true iris-like texture. The dot-matrix lenses also allow some of the natural iris color to be seen through the holes of the matrix pattern [21]. Figures 7(a) and 7(b) are examples of these two subjects. We note that the lenses seen in Figure 7(a) were sample cosmetic lenses provided to the subject and have the letters "DEMO" written on the lens. Additionally, the true iris texture is more visible in the area around these letters.

While cosmetic contact lenses are known to disrupt iris recognition and techniques have been proposed to detect dot-matrix style lenses, we investigate their effect by considering two sets of matches for each subject and variation of cosmetic lens: 1) we consider matches between two images where the subject is wearing cosmetic lenses in both images and 2) we consider matches between images where in one image the subject is wearing the cosmetic lens and in the other image the subject is not wearing the lens. We also distinguish between matches of the first variety between images acquired at the same session and images acquired at different sessions.

IV. EXPERIMENTAL RESULTS FROM THE IRISBEE SYSTEM WITH ACTIVE CONTOURS

We separated the match scores for the IrisBEE results with active contours into groups based on the categorization of the probe and gallery images in the match. We then computed the mean match score and false reject rate for each group. Table I shows these means and false reject rates at a 0.32 fractional Hamming distance threshold. An entry in the table shows the mean match score for matches between images of the two types in the corresponding row and column. We note that the symmetric nature of the Hamming distance allows us to treat, for example, a gallery image of Category 1 and a probe image of Category 2 as equivalent to a gallery image of Category 2 and a probe image of Category 1.

Based upon our previous results [23], we expected the "Category 0 versus Category 0" (no contact lenses) matches to outperform the other categories. We find that the Category 0 matches have a match Hamming distance mean of 0.131 and a false reject rate of 1.17%. This is the best performance among all categories.

We observe an increase in mean match Hamming distance with each increasing contact degree complexity. While Category 1 matches have an increase over Category 0 matches in mean match Hamming distance at 0.156, they show show little degradation in performance with a false reject rate of 1.67% compared to 1.17% for Category 0 matches. The Category 2 matches have a mean

(a) Subject with demo dot-matrix patter cosmetic lens. Image name: 05303d782.

(b) Subject with dot-matrix pattern cosmetic lens. Image name: 06244d20.

Fig. 7. Example of Dot-matrix pattern cosmetic lenses. We observe the letters "DEMO" written on the lens shown in Figure 7(a).

TABLE I

MATCH HAMMING DISTANCE MEANS, FALSE REJECT RATES (AT A MATCH SCORE THRESHOLD OF 0.32) AND THE NUMBER
OF MATCHES FOR COMPARISONS BETWEEN DIFFERENT CATEGORIES OF CONTACT LENS SUBJECTS FOR THE IRISBEE
SYSTEM WITH ACTIVE CONTOURS. THE COMPLEXITY OF CONTACT LENSES INVOLVED INCREASES MOVING DOWN AND
ACROSS THE TABLE.

	Category 0	Category 1	Category 2	Category 3	Category 4
Category 0					
Mean	**0.1314**	0.1675	0.1791	0.1813	0.1591
FRR	**1.17%**	5.66%	6.02%	2.67%	0.49%
Number	**447875**	34459	15760	3000	412
Category 1					
Mean		**0.1559**	0.1544	0.2103	NA
FRR		**1.67%**	1.0%	2.57%	NA
Number		**311719**	1114	740	0
Category 2					
Mean			**0.1760**	0.1529	NA
FRR			**3.72%**	9.0%	NA
Number			**211687**	32	0
Category 3					
Mean				**0.1803**	NA
FRR				**5.77%**	NA
Number				**86634**	0
Category 4					
Mean					**0.2824**
FRR					**40.10%**
Number					**42329**

match Hamming distance of 0.176 and a false reject rate of 3.72%, which is an approximate three-fold increase over the false reject rate for Category 0 matches.

The trend continues with the Category 3 matches, which have a false reject rate of 5.77%, a 5-fold increase from the 1.17% for the Category 0 matches.

Category 4 suffers the worst performance with a mean match Hamming distance of 0.282 for correct matches and a false reject rate of 40.1% which is almost 35 times that of Category 0 matches. We observe that using this version of the IrisBEE system, a correct match is almost as likely to be falsely rejected as it is to be accepted.

Match Distributions for Levels of Contacts
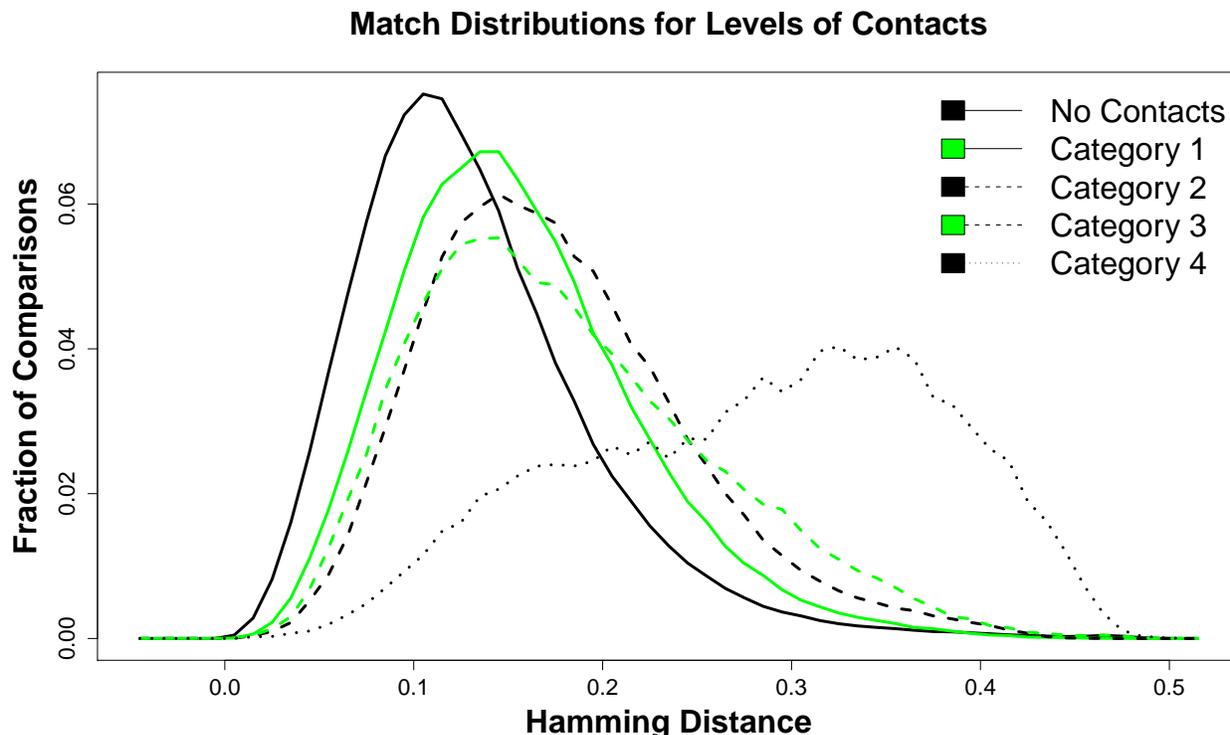
Fig. 8. Correct Match distributions for categories of contacts from the IrisBEE system with active contours. Lower scores are better matches.

Figure 8 exhibits the "authentic" or "true match" distributions for each of the categories from the IrisBEE system with active contours. Category 4 matches are clearly much worse than the other categories, and Category 3 has a significant number of matches with very high Hamming distances. In general we observe a shift to the right for the correct match distributions of Categories 2, 3, and 4.

Across all contacts we found the mean Hamming distance was 0.174 for matches between images of the same type of contact lenses with a false reject rate of 5.38%. This represents an approximate five-fold increase in false reject rate over the 1.17% false reject rate for Category 0 images. Figure 9 shows the score distributions for correct Category 0 matches and matches between images of the same contact lens category. The contact lens distribution is shifted to the right (degraded match scores) from the no contact lens distribution. The significant tail we observe in the contact lens match score distribution is a result of the degradation for Category
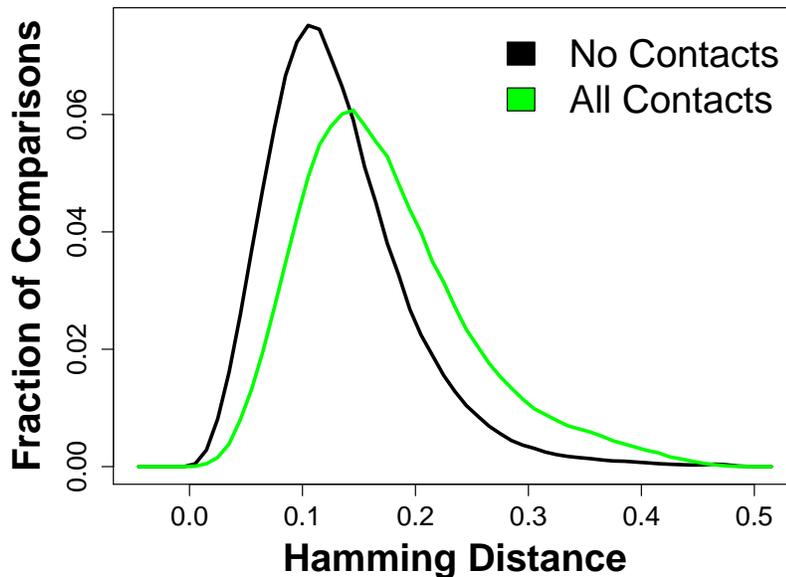
Fig. 9. True match score distributions for all contacts (Categories 1, 2, 3, and 4 combined) and no contacts from the IrisBEE system with active contours. Lower scores are better matches.

TABLE II

MEAN HAMMING DISTANCE AND FALSE ACCEPT RATES FOR COMPARISONS OF IMAGES OF THE SAME TYPE OF CONTACT LENS BUT DIFFERENT IRISES USING THE IRISBEE SYSTEM WITH ACTIVE CONTOURS.

	No Contacts	Category 1	Category 2	Category 3	Category 4
HD	0.452	0.452	0.453	0.454	0.453
FAR	3.19 E-6	9.99 E-7	1.59E-7	2.60E-6	0

3 and Category 4 matches observed earlier.

For non matches between images of the same type, we computed the mean Hamming distance and the false accept rate at a 0.32 Hamming distance threshold. These results are shown in Table II. As expected we do not observe a large shift in the non match Hamming distance means or false accept rates.

Figure 10 shows the distributions of the non matches for each of the categories. These distributions are very similar for all categories, indicating that the non match distributions are
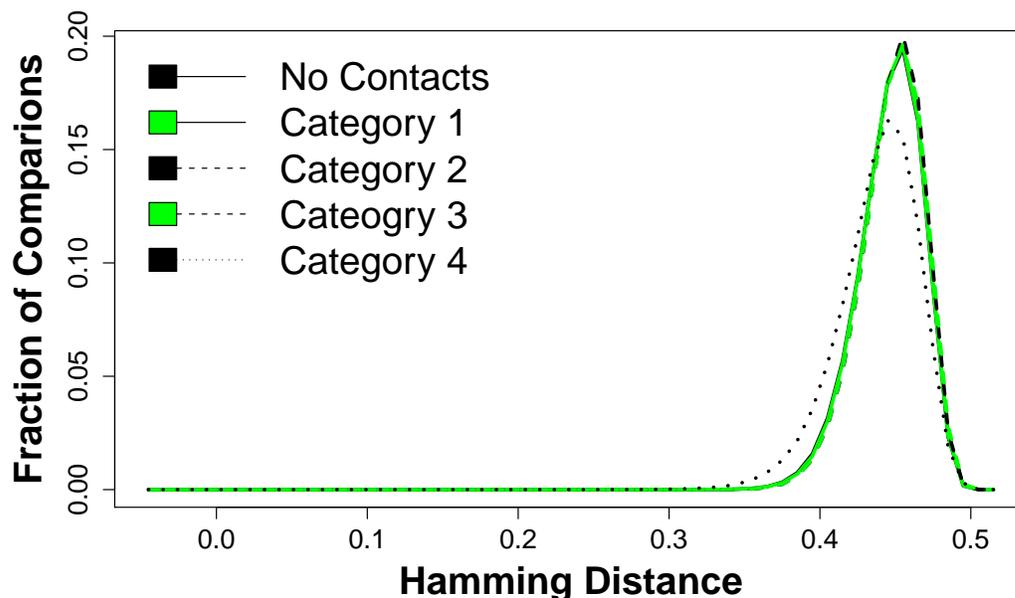
Fig. 10. False match score distributions for all contacts (Categories 0, 1, 2, 3, and 4) from the IrisBEE system with active contours.

not affected by contact lenses.

The ROC curves for all contacts and no contacts for the IrisBEE system with active contours are shown in figure 11. We observe that the contacts have a marked degradation in performance compared to no contacts.

V. EXPERIMENTAL RESULTS FROM THE VERIEYE SYSTEM

We repeated our experiments using the commercial VeriEye Iris SDK (version 2.2) from Neurotechnology [13]. The match score reported by VeriEye is not a fractional Hamming distance. The system reports a match score between zero and 3235, where 3235 is reported for identical iris images and zero is reported if the two images are determined to represent different irises. The average match score for two different images of the same iris is about 400. However, unlike the IrisBEE system, the scores are not symmetric and the score reported for Image 1 vs Image 2 is close to but not always the same as the score reported for Image 2 vs Image 1. Therefore, we computed the average of these two scores and used this average in
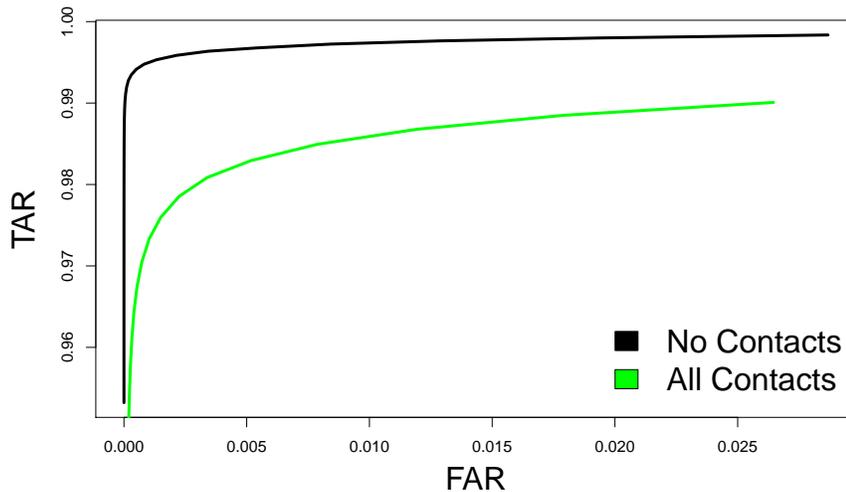
Fig. 11. ROC curve for all contacts (Categories 0, 1, 2, 3, and 4) and no contacts from the IrisBEE system with active contours.

reporting our results.

The resulting match score means and false reject rates for the different types of matches are presented in Table III. We used 43 as a threshold for computing the false reject rate because at this threshold the false accept rate for Category 0 was the same as the false accept rate for Category 0 using the IrisBEE system. We observe that the mean match score degrades as the artifact left by the contact lens increases.

The Category 1 and Category 2 contact matches have mean match scores lower than that of the Category 0 matches, but their false reject rates are only slightly worse than that of Category 0. We observe that matches between images containing the Category 3 and Category 4 contact lenses suffer significant degradation in quality. The false reject rate of 0.70% for Category 3 contacts is about eighteen times that of "no contact" matches, and the false reject rate of 2.02% for Category 4 (or gas-permeable) contact lens matches is approximately fifty times that of "no contact" matches. While the false reject rates are in general much lower than those for IrisBEE, the same general trends and percentage of degradation are observed.

Figure 12 shows the score distributions of the different categories of true matches. Again, we observe a clear shift in the match distribution for the categories of contacts, especially on the

Match Distributions for Categories of Contacts
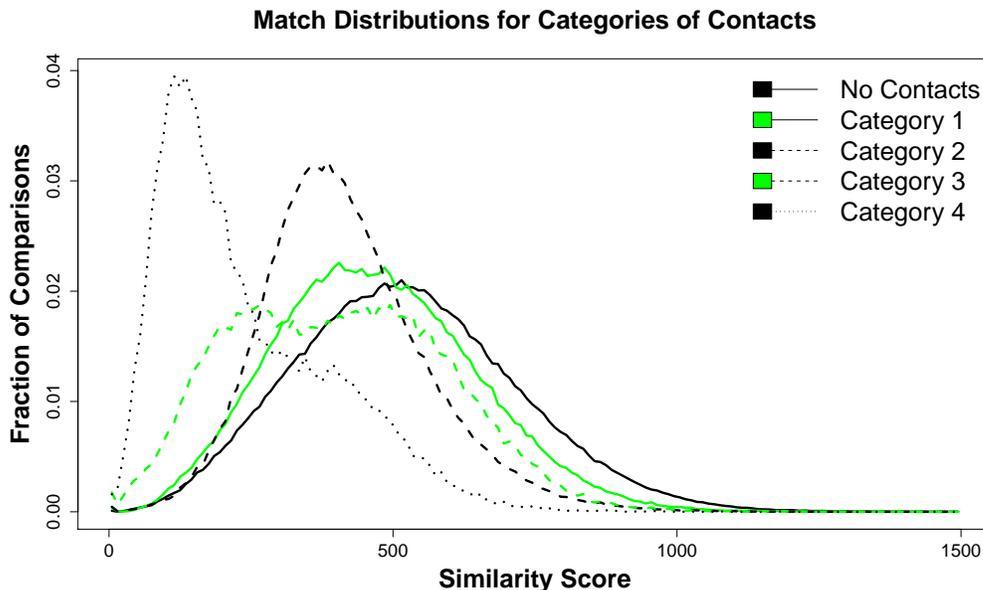
Fig. 12. True match score distributions for categories of contacts from the VeriEye system. Higher scores are better matches.

right tail. We note that the right tail of the distributions represents matches of particularly high quality. The Category 0 match distribution indicates that high quality matches are more likely to occur with the Category 0 images than any of the other categories. However, the left tail of Category 1 and Category 2 (the false reject end of the distributions) are much closer to the left tail of Category 0. Categories 3 and 4, however, do show significant degradation compared to Category 0 throughout their distributions.

Across all matches between images containing the same type of contact lenses, we found the average match score was 428.5 as compared to the average match score of 523.5 for Category 0 images. Figure 13 illustrates the true match score distributions for Category 0 and matches between images of the same contact lens category. We observe a shift in the distribution for contacts indicating a degradation in match quality for these images. The false reject rate for all contacts was 0.28% compared to 0.04% for no contacts, a seven-fold increase.

Figure 14 shows the distributions of the non matches for each of the categories. These distributions are almost exactly the same. The mean non match scores for all categories were all about 1.3 as the majority of non matches result in a score of 0.

We plotted the ROC curves for all contacts and no contacts for the VeriEye system and show
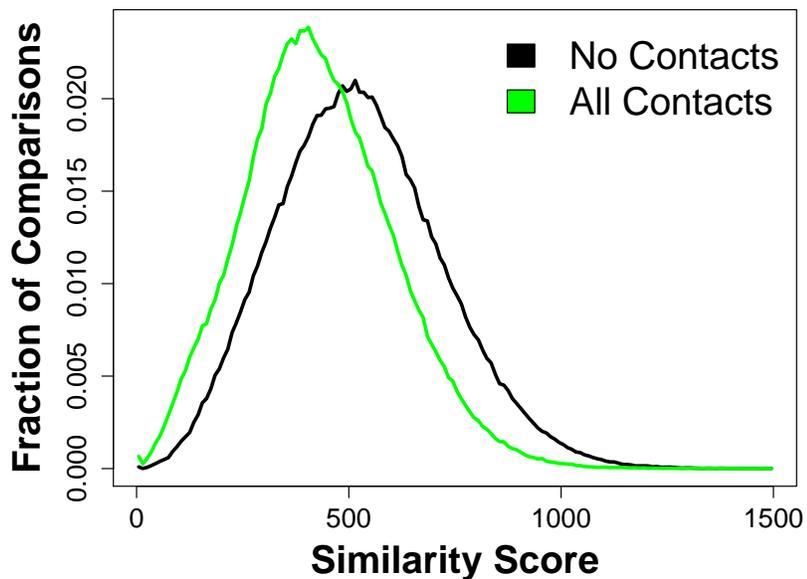
Fig. 13. True match score distributions for all contacts and no contacts from the VeriEye system. Higher scores are better matches.
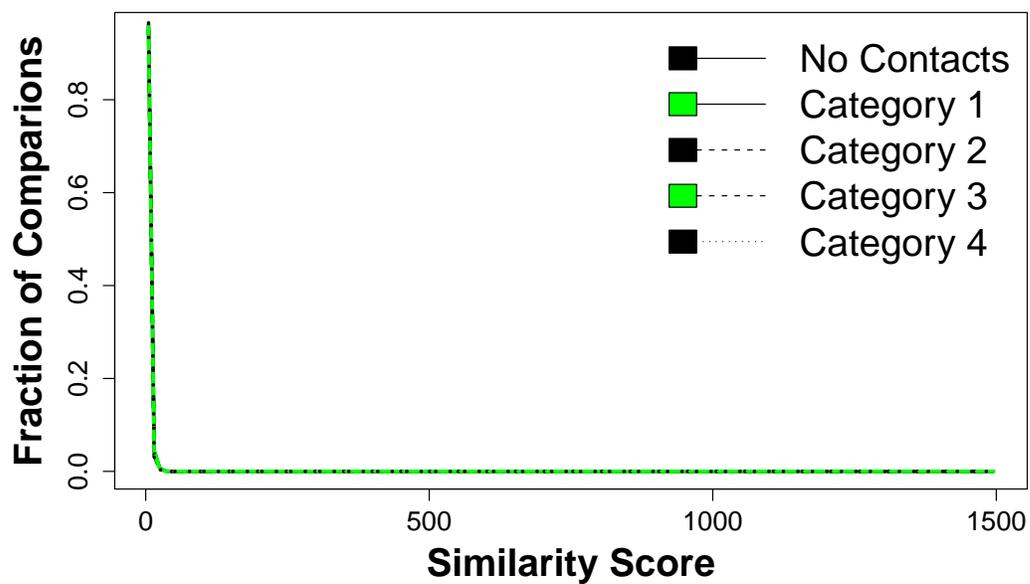
Fig. 14. False match score distributions for all contacts (Categories 0, 1, 2, 3, and 4) from the VeriEye system.

TABLE III

VERIEYE MATCH SCORE MEANS, FALSE REJECT RATES AND THE NUMBER OF MATCHES FOR COMPARISONS BETWEEN

DIFFERENT CATEGORIES OF CONTACT LENS SUBJECTS USING ALL IMAGES.

		Category 0	Category 1	Category 2	Category 3	Category 4
Category 0						
	Mean	**523.5**	512.8	470.5	271.5	458.5
	FRR	**0.04%**	0.44%	3.20%	0.0%	0%
	Number	**447612**	34471	15801	3000	412
Category 1						
	Mean		**469.9**	448.0	214.9	NA
	FRR		**0.05%**	0.0%	0.0%	NA
	Number		**311908**	1114	740	0
Category 2						
	Mean			**415.7**	872.2	NA
	FRR			**0.08%**	0%	NA
	Number			**211747**	32	0
Category 3						
	Mean				**403.2**	NA
	FRR				**0.70%**	NA
	Number				**86649**	0
Category 4						
	Mean					**239.8**
	FRR					**2.02%**
	Number					**42442**

this in figure 15. We observe that the performance of the contacts categories is much worse than that of the no contacts category.

The VeriEye system performs better overall than the IrisBEE system. Based upon the performance observed, it seems likely that the VeriEye system has some built in method of locating and accounting for contact lenses, especially the gas-permeable lenses. However, degradation due to contact lenses is still present in the VeriEye system.

VI. EXPERIMENTAL RESULTS FROM THE CMU SYSTEM

We repeated our experiments using a system provided by Carnegie Mellon University [28]. The match score reported by this system is also a fractional Hamming distance.
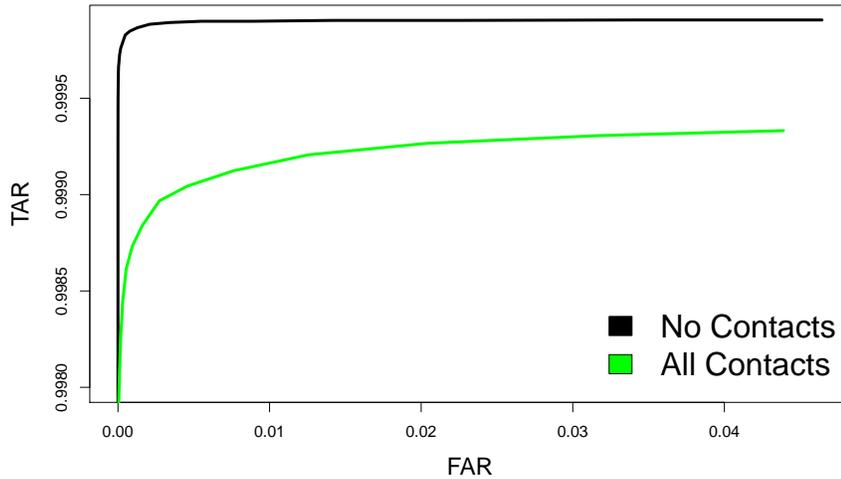
Fig. 15.   ROC curve for all contacts (Categories 0, 1, 2, 3, and 4) and no contacts from the VeriEye system.

The resulting match score means and false reject rates for the different types of matches are presented in Table IV. We used 0.28 as a threshold for computing the false reject rate because at this threshold the false accept rate for Category 0 was the same as the false accept rate for Category 0 using the IrisBEE system. We observe that the mean match score degrades as the artifact left by the contact lens increases.

The Category 1 and Category 2 contact matches have mean match scores lower than that of the Category 0 matches, but the false reject rate for Category 2 is about the same as the false reject rate for Category 0. The false reject rate for Category 1 is about a 75% increase from Category 0.

We observe that matches between images containing the Category 3 and Category 4 contact lenses suffer significant degradation in quality. The false reject rate of 9.3% for Category 3 contacts is about four times that of "no contact" matches, and the false reject rate of 47.5% for Category 4 (or gas-permeable) contact lens matches is approximately twenty times that of "no contact" matches.

Figure 16 shows the score distributions of the different categories of true matches. Again, we observe a clear shift in the match distribution for the categories of contacts. At the far right

TABLE IV

CMU MATCH SCORE MEANS, FALSE REJECT RATES AND THE NUMBER OF MATCHES FOR COMPARISONS BETWEEN

DIFFERENT CATEGORIES OF CONTACT LENS SUBJECTS USING ALL IMAGES.

		Category 0	Category 1	Category 2	Category 3	Category 4
Category 0						
	Mean	**0.153**	0.167	0.166	0.169	0.196
	FRR	**2.23%**	2.64%	1.69%	0.5%	1.70%
	Number	**442123**	34471	15801	3000	412
Category 1						
	Mean		**0.164**	0.131	0.258	NA
	FRR		**3.89%**	3.32%	27.7%	NA
	Number		**311908**	1114	740	0
Category 2						
	Mean			**0.166**	NA	NA
	FRR			**2.17%**	0%	NA
	Number			**211747**	32	0
Category 3						
	Mean				**0.184**	NA
	FRR				**9.28%**	NA
	Number				**86604**	0
Category 4						
	Mean					**0.273**
	FRR					**47.45%**
	Number					**42442**

tails of the distributions, we observe that Categories 0, 1, and 2 are very similar, but there is a significant increase in the right tails of Categories 3 and 4.

Across all matches between images containing the same type of contact lenses, we found the average match score was 0.174 as compared to the average match score of 0.153 for Category 0 images. The false reject rate of all contacts was 6.88% compared to 2.23% for Category 0, which is an approximate three-fold increase. Figure 17 illustrates the true match score distributions for Category 0 and matches between images of the same contact lens category. We observe a shift in the distribution for contacts indicating a degradation in match quality for these images.

Figure 18 shows the distributions of the non matches for each of the categories. Again, as in the IrisBEE and the VeriEye systems, these distributions are almost exactly the same. The mean

Match Distributions for Categories of Contacts
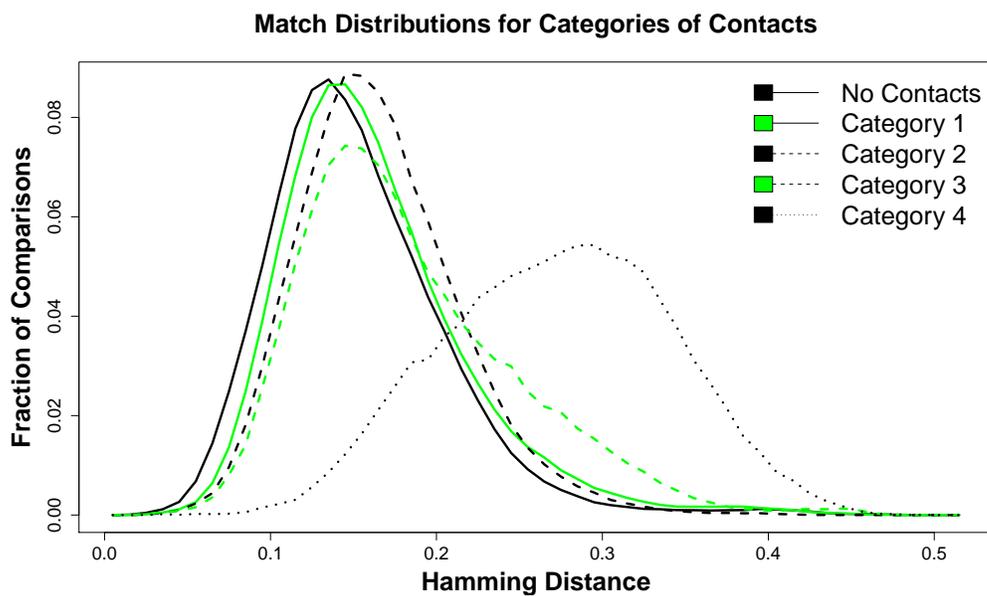
Fig. 16.   True match score distributions for categories of contacts from the CMU system. Lower scores are better matches.
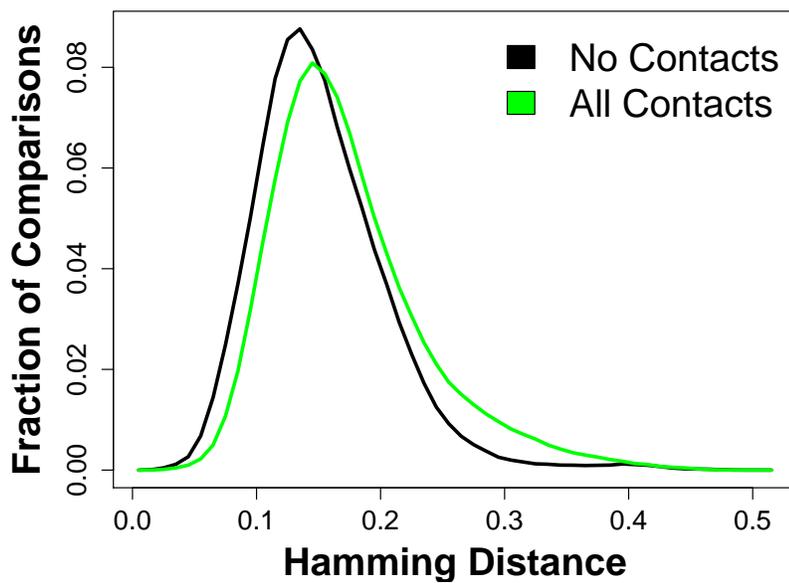
Fig. 17.   True match score distributions for all contacts and no contacts from the CMU system. Lower scores are better matches.
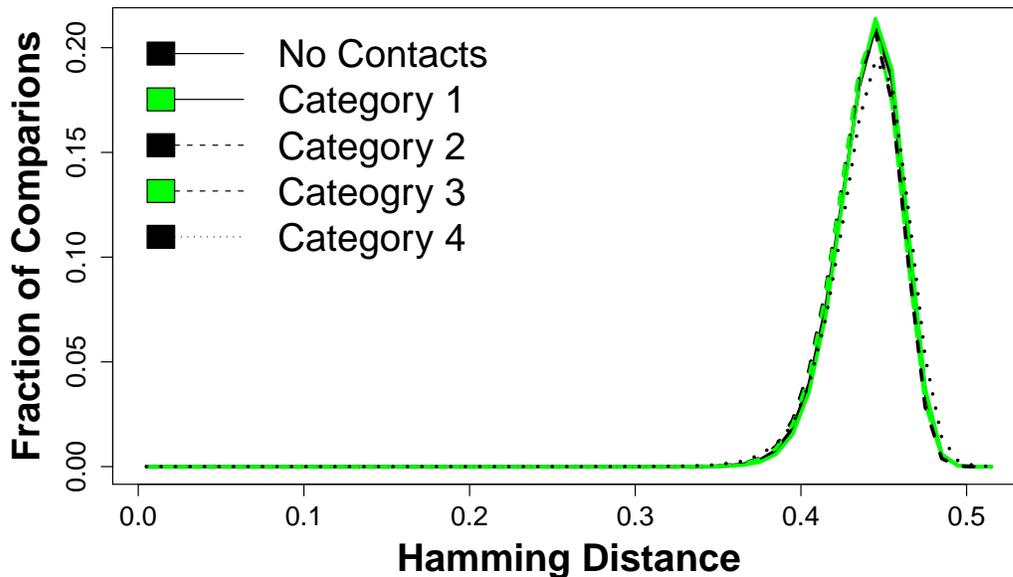
Fig. 18.   False match score distributions for all contacts (Categories 0, 1, 2, 3, and 4) from the CMU system.

non match scores for all categories were 0.44; the false accept rate was about 1 in a million for all categories at a 0.28 threshold.

We plotted the ROC curves for all contacts and no contacts for the CMU system and show this in figure 19. As in the other systems, we observe that the performance of the contacts categories is much worse than that of the no contacts category.

VII. COSMETIC CONTACT LENS RESULTS

We only have three subjects who wore cosmetic lenses, and each subject has a different type of lens as described in section III-D. Two of the subjects only participated in one acquisition session with and without their cosmetic lenses. One subject was a frequent participant in acquisitions, so we consider both same session matches and different session matches for this subject. We present the results for matches involving these subjects in Table V for the IrisBEE system and in Table VI for the VeriEye system.

Across all types of cosmetic lenses, we observe that two cosmetic contact lens images from the same iris acquired at the same session can match as a true accept. Subject 06244 who was
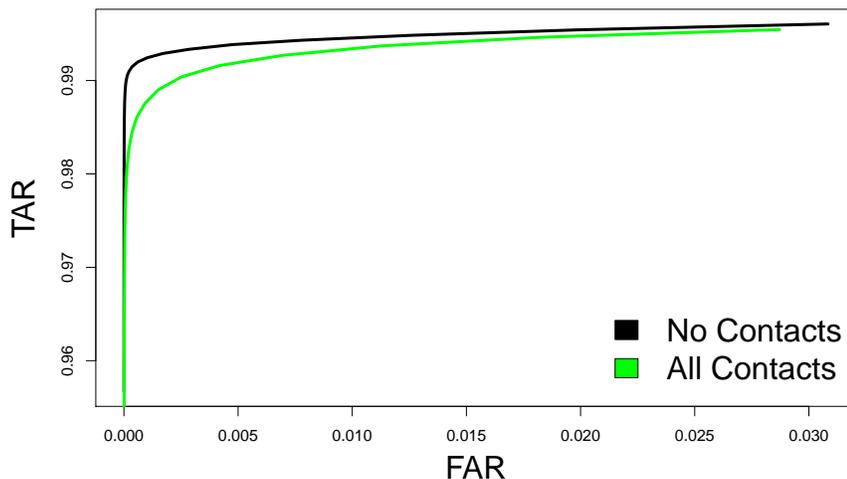
Fig. 19.   ROC curve for all contacts (Categories 0, 1, 2, 3, and 4) and no contacts from the CMU system.

TABLE V

IRISBEE MATCH SCORE, FALSE REJECT RATES AND NUMBER OF MATCHES FOR EACH OF THE THREE COSMETIC LENS

SUBJECTS.

Subject	Type of Match	Match Mean	Number Matches	FRR
04780 - same session	Cosmetic vs Cosmetic	0.2972	22	54.5%
04780 - different session	Cosmetic vs Cosmetic	0.4392	84	100%
Pad-printed lenses	Cosmetic vs not Cosmetic	0.4020	240	95.4%
05303 - same session	Cosmetic vs Cosmetic	0.3729	18	83.3%
Demo dot-matrix lenses	Cosmetic vs not Cosmetic	0.4003	39	94.9%
06244 - same session	Cosmetic vs Cosmetic	0.2164	25	0%
Dot-matrix lenses	Cosmetic vs not Cosmetic	0.4085	66	100%

wearing regular dot-matrix lenses experienced a 0% false reject rate for comparisons between two cosmetic contact lens images. Subject 05303 suffered significant degradation for the same type of comparisons, but this may be due to the added artifact of the "DEMO" written on the lenses. The pad-printed lens wearer, subject 04780, experienced a false reject rate of 54.5%. While it is possible to produce images that will match well at the same session, the results show that this is certainly not guaranteed.

TABLE VI

VeriEye match score, false reject rates and number of matches for each of the three cosmetic lens
subjects.

Subject	Type of Match	Match Mean	Number Matches	FRR
04780 - same session	Cosmetic vs Cosmetic	267.84	22	4.5%
04780 - different session	Cosmetic vs Cosmetic	4.65	84	100%
Pad-printed lenses	Cosmetic vs not Cosmetic	16.3	240	99.6%
05303	Cosmetic vs Cosmetic	160.0	18	61.1%
Demo dot-matrix lenses	Cosmetic vs not Cosmetic	28.1	39	94.9%
06244	Cosmetic vs Cosmetic	573.9	25	0%
Dot-matrix lenses	Cosmetic vs not Cosmetic	7.0	66	100%

The only returning cosmetic lens subject had a 100% false reject rate for matches between cosmetic lens images from different sessions. However, this subject wears pad-printed lenses that are not uniform throughout and are specifically designed to imitate an iris texture. Within an acquisition session, the contact lens is unlikely to experience significant rotation so the overall texture of the iris and the lens will remain relatively constant, thus good match scores are possible. On the other hand, because the contact lens has no preferred orientation, the texture of the lens is likely to align with the true iris in a different manner in different images. Therefore, the combined texture of the lens and the true iris will cause images from different sessions to not match well.

We do not have data from different sessions for the dot-matrix lens subjects. This would be an important extension of these results in the future. The dot-matrix lenses are designed such that they have the same general texture throughout. Perhaps the grid-like texture of these lenses will cause matches for images from different sessions to produce good match scores.

We also considered the possibility that a subject could masquerade as another subject by wearing cosmetic contact lenses. To this end, we computed the false accept rates for both the IrisBEE and the VeriEye algorithms. Across all possible matches, the false accept rate was 0% so that even matches between images of one subject's right and left eye from the same session both wearing the same type of cosmetic lens resulted in a match score that was a true reject. This result suggests that using these types of cosmetic lenses a false accept is unlikely to occur.
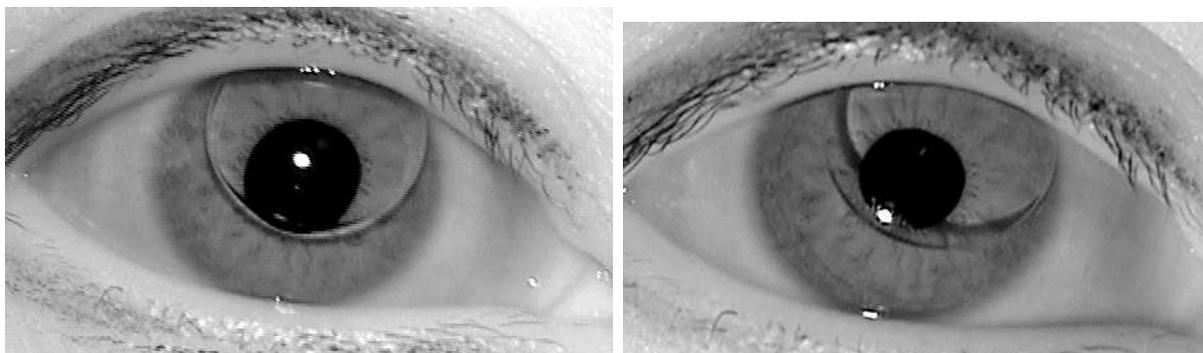
(a) Gallery Image. Image name: 04593d617.   (b) Probe Image. Image name: 04593d624.
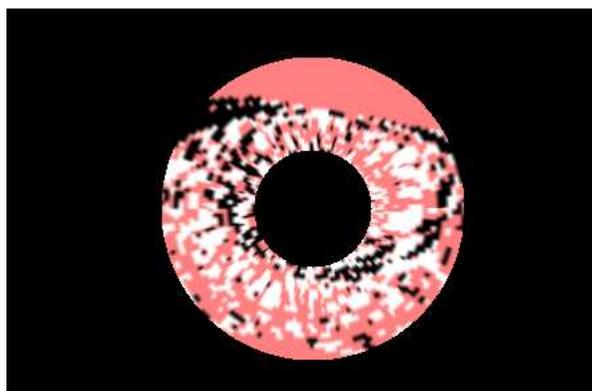
(c) Unmatching bits in comparison

Fig. 20. The black regions represent bits that did not agree in the match between Image 04593d617 and image 04593d624. The red regions are bits that were masked in one or both of the images during the segmentation and the white regions represent bits that did agree in the match. With the exception of some eyelid occlusions, the black regions are around the contact lens boundary. This pair of images resulted in a Hamming distance of 0.315.

However, these lenses are specifically made to allow for areas of the true iris to be visible through gaps in the colored portion of the lens. If a lens was created so that the entire lens was colored with a fake iris texture, a false accept may be possible.

VIII. VISIBLE EFFECTS OF CONTACT LENSES

In this section we present our work to more precisely pinpoint the effect of certain contact lenses. We compare the aligned iris codes for two images from the same iris. Each bit in the iris code corresponds to a location in the original iris image; thus determining bits that do not match in the two iris codes allows us to locate the areas of the iris texture that disrupt recognition. We mark the regions that were masked by one or both of the images as red, regions with bits that
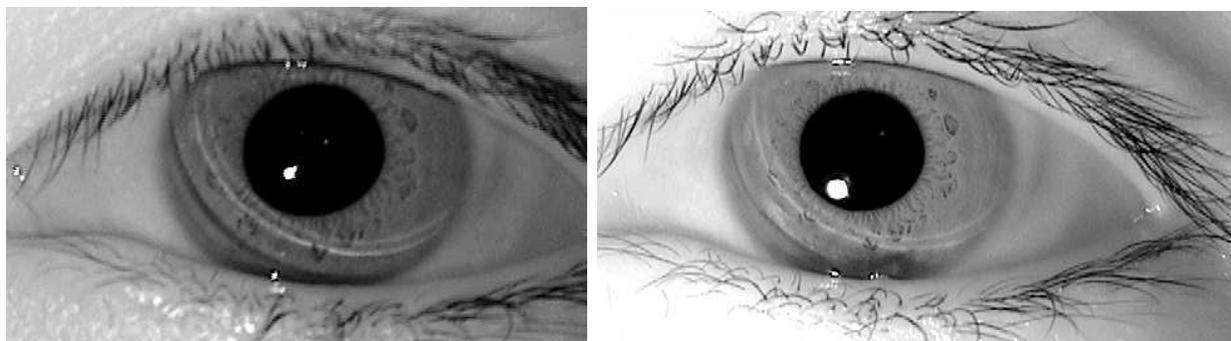
(a) Gallery Image. Image name: 04221d1067.

(b) Probe Image. Image name: 04221d1070.
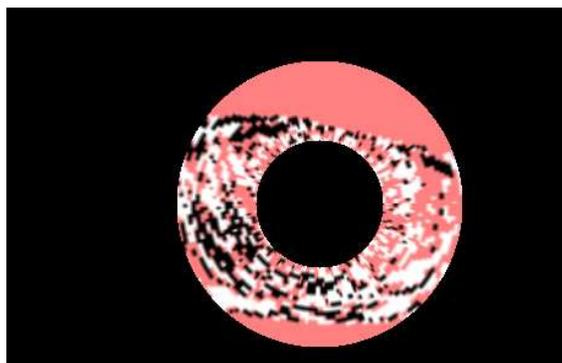
(c) Unmatching bits in comparison

Fig. 21. The black regions represent bits that did not agree in the match between image 04221d1067 and image 04221d1070. The red regions are bits that were masked in one or both of the images during the segmentation and the white regions represent bits that did agree in the comparison. With the exception of some eyelid occlusions, the black regions mostly are found near the artifact caused by the contact lens. This pair of images resulted in a Hamming distance of 0.364.

matched in the two iris codes as white, and regions with bits that did not match in the two iris codes as black. We present the results of this for two subjects.

Subject 04593 wears gas-permeable contacts and as Figure 20 shows, the gas-permeable contacts mostly affect the iris code along its boundary. We do note that we have some eyelid occlusion causing some bits to not match in those regions as well. The region of the iris covered by the lens is mostly unaffected.

Subject 04221 had ill-fitting contact lenses that left a substantial artifact on his/her iris. As Figure 21 shows, this artifact significantly affected the match score. Again, we observe some bits not matching due to eyelid occlusion, but the regions of the contact lens' artifact are largely affected.

In both of these figures, the images used for comparison were acquired in the same acquisition session within seconds of each other. Comparing images from different acquisition sessions shows an even more noticeable effect as the contact lens (and its resulting artifact) will not be in the same general location as they are in these pairings.

IX. Discussion and Future Work

The results presented here show that iris match quality degrades as the artifact left on the iris from the contact lens increases. Even matches between images with no visible artifacts or with only a circular ring visible on the iris experience match quality degradation for the VeriEye, CMU, and IrisBEE with active contours systems. While the performance of VeriEye is better across all categories than the IrisBEE and CMU systems, the contact effect is apparent in all three.

Matches involving images of soft contact lenses with letter indicators or ill-fitting artifacts experience a false reject rate about five times that of matches between images of no contacts on the IrisBEE system with active contours, 18 times for the VeriEye system, and about three times for the CMU system. Gas-permeable contact lenses cause a 35-fold increase for IrisBEE with active contours, a 50-fold increase on the VeriEye system, and a 20-fold increase on the CMU system. We speculate that this degradation is primarily due to the distortion from the lens not being in the same place on two iris images acquired at two different times. For all algorithms considered, the false accept rate is not be affected by contact lenses, and the non match distributions for all categories are nearly identical.

In a traditional verification scenario, the security of the system is unlikely to be affected by the presence of contact lenses. Problems mainly arise with the ease-of-use of a system. However, security will be an issue in a scenario such as a "watch list" where a system is created to deny access to a list of people while admitting all those not matching the people on the list. In this situation, a person who should be denied access could wear gas-permeable contact lenses which at this time are not detected by the system, and they are 20-50 times more likely to evade detection on the watch list.

Our results from comparing the Category 0 images to the Categories 1-4 images show that the false reject rates for these cross-category comparisons are generally lower than the rates for comparing contact to contact images. One possibility to mitigate this effect is to require all

(a) Gallery Image. Image name: 04622d615.
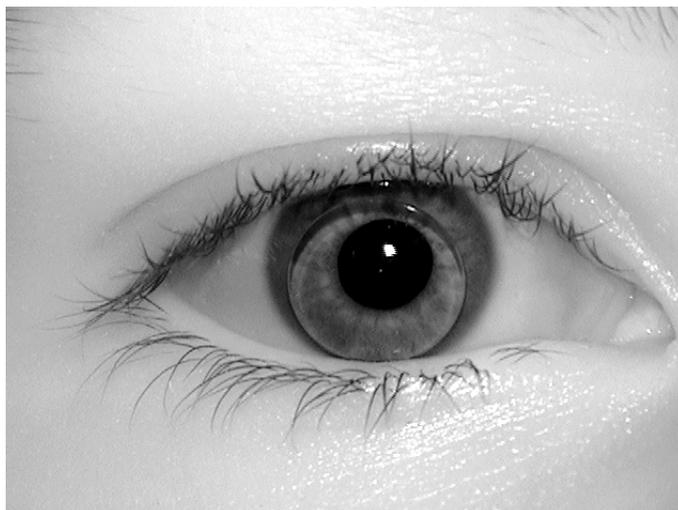
(b) Probe Image. Image name: 04622d673.

Fig. 22. These two images matched in the VeriEye system with a match score of 335, but resulted in a Hamming distance of 0.407 in the IrisBEE system with active contours. The high (good) match score in VeriEye indicates that this system accounts for the gas-permeable contact lens in a way that does not occur in the IrisBEE system resulting in the high (bad) Hamming distance.

subjects to enroll without contact lenses. Then at the time of recognition or verification, the subject can wear or not wear lenses. It may also be useful to allow the subject to enroll with two images - one with contact lenses and one without. In a recognition or verification scenario, the subject can be compared to both images.

We have also observed that in general VeriEye outperforms IrisBEE and CMU. It is possible that this system has some methodology of checking for contact lenses or that the algorithm used inherently reduces the contact lens effect. The false reject rate for Category 4 matches was about 40% in IrisBEE and 48% in CMU, but only about 2% in VeriEye. For example, Figure 22 shows two images that matched poorly in IrisBEE but VeriEye was able to match within the true accept range. As we have shown in Section VIII, the boundary regions of the gas-permeable contact lenses specifically disrupted the IrisBEE matcher. One way to improve the effect of contact lenses on the IrisBEE software is to implement a method to locate these boundaries and mask them appropriately. It may also be possible to use video to locate and track a contact lens as it moves from frame to frame.

Developing techniques to detect and mitigate the contact lens effect is important and necessary in order to attain a successful iris biometric recognition system. Daugman [12] and Wei et al. [4] present methods to detect cosmetic lenses, but our results show that non-cosmetic lenses need to be detected as well.

We have shown that images of a person wearing cosmetic lenses from the same session will result in possible true accept match scores, whereas images from different sessions never match well for both the VeriEye and the IrisBEE algorithms. If we are unable to obtain a high quality match score for the same iris across different sessions when the subject wears cosmetic lenses, it is unlikely a person could use a cosmetic lens to successfully masquerade as someone else. The cosmetic lenses in this study are all specifically designed with holes in the color and texture pattern in order to allow the true iris to be seen. If a lens was created so that the entire lens is colored and textured, it may be possible for a person to assume another's identity. Of particular interest would be if a person could replicate another's iris texture on a cosmetic lens and pose as that person.

As noted, we have very limited data for the cosmetic lenses and no data for different sessions for dot-matrix lens. To make any further conclusions, it would be necessary to repeat this experiment on more subjects and more data.

Our data set consists of the first approximately 200 subjects participating in acquisitions in the 2004-2005 academic year. The subjects are primarily a sampling of Midwestern college aged students. It is likely that a sampling of contact lens wearers today would reveal even more unusual patterns than noted here. As the contact lens industry continues to evolve, the

complexity of available lenses increases. For example, lenses are increasingly manufactured with letters, numbers or other markings. Also, bifocal gas permeable contact lenses [26] and multifocal lenses are becoming more popular. None of the subjects in our data set appear to be wearing these types of lenses, but they would likely cause a significant disturbance in iris recognition. It is important that the effects of contact lenses be addressed, especially as the variety of available lenses grows.

X. ACKNOWLEDGEMENT

REFERENCES

[1] J. Daugman, "How Iris Recognition Works," *IEEE Transactions On Circuits and Systems for Video Technology,* 14(1):21-30, 2004.

[2] LG. http://www.lgiris.com/, accessed April 2009.

[3] X. Liu, K. W. Bowyer, P. J. Flynn. "Experiments with an improved iris segmentation algorithm," *Fourth IEEE Workshop on Automatic Identification Technologies*", 118-123, Oct 2005.

[4] Z. Wei, X. Qui, Z. Sun, and T. Tan, "Counterfeit Iris Detection Based on Texture Analysis," *Proc. of IEEE Int'l Conf. on Pattern Recognition* 2008.

[5] J. Ali, A. Hassanien,"An Iris Recognition System to Enhance E-security Environment Based on Wavelet Theory," *Advanced Modeling and Optimization* 5(2): 2003.

[6] G. Williams, "Iris Recognition Technology," *IEEE Aerospace and Electronic Systems*, 12(4): 1997.

[7] M. Negin, T. Chmielewski Jr., M. Salganicoff, T. Camus, U. von Seelen, P. Venetianer, G. Zhang, "An Iris Biometric System for Public and Personal Use," *Computer* 33(2): 2000.

[8] X. Liu. "Optimizations in Iris Recognition." PhD Dissertation, University of Notre Dame, 2006.

[9] P. J. Philips, K. W. Bowyer, P. J. Flynn, X. Liu, W. T. Scruggs, "The Iris Challenge Evaluation 2005," *Proc. of 2008 IEEE Conference on Biometrics: Theory, Applications, and Systems.* Sept 2008.

[10] P. J. Phillips, W. T. Scruggs, A. O'Toole, P. J. Flynn, K. W. Bowyer, C. L. Schott, M. Sharpe, "FRVT 2006 and ICE 2006 Large-Scale Experimental Results." *IEEE Transactions on Systems, Man, and Cybernetics.*, in press.

[11] K. Hollingsworth, K. W. Bowyer, P. J. Flynn, "The Best Bits in an Iris Code," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 31(6): 964-973, 2009.

[12] J. Daugman, "Demodulation by Complex-Valued Wavelets for Stochastic Pattern Recognition," *Int'l Journal of Wavelets Multiresolution and Information Processing* 1(1): 1-17, March 2003.

[13] VeriEye Iris Recognition Technology. http://www.neurotechnology.com/verieye.html, accessed November 2008.

[14] http://en.wikipedia.org/wiki/Contact_lens, accessed May 2009.

[15] J. Daugman, "New Methods in Iris Recognition," *IEEE Transactions On Systems, Man, and Cybernetics.* 37(5):1167-1175, Oct 2007.

[16] Acuvue Contact Lens Products . http://www.acuvue.com/products-home.htm, accessed May 2009.

[17] J. Daugman "The Importance of Being random: Statistical Principles of Iris Recognition," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(2): 279-291, 2003.

[18] 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.

[19] S. Ring and K. W. Bowyer, "Detection of Iris Texture Distortions by Analyzing Iris Code Matching Results" *Proc. of 2008 IEEE Conference on Biometrics: Theory, Applications, and Systems.* Sept. 2008.

[20] http://www.daviddarling.info/encyclopedia/T/toric_contact_lenses.html, accessed June 2009.

[21] M. Horn and A. Bruce, "Manual of Contact Lens Prescribing and Fitting," Elsevier Health Sciences, 687-691, 2006.

[22] D.G. Streibig, "Colored Contact Lens," US Patent 7,296,891 B2, 2007.

[23] S. E. Baker, K. W. Bowyer, and P. J. Flynn. "Contact Lenses: Handle with Care for Iris Recognition." *Proc. of 2009 IEEE Conference on Biometrics: Theory, Applications and Systems* in press.

[24] K.W. Bowyer and P.J. Flynn, "The ND-IRIS-0405 Iris Image Dataset." Notre Dame CVRL Technical Report, http://www.nd.edu/ cvrl/papers/ND-IRIS-0405.pdf.

[25] Gas Permeable (GP) Contact Lenses. http://www.allaboutvision.com/contacts/rgps.htm, accessed August 2009.

[26] Bifocal and Multifocal Contact Lenses. http://www.allaboutvision.com/contacts/bifocal.htm, accessed August 2009.

[27] T. Peters, "Effects of Segmentation Routine and Acquisition Environment on Iris Recognition." M.S. Thesis, University of Notre Dame, 2009.

[28] J. Thornton, M. Savvides, B.V.K. Vijaya Kumar, "A Bayesian Approach to Deformed Pattern Matching of Iris Images," *IEEE Transactions on Pattern Analysis and Machine Intelligence,* 29(4): 596-606, April 2007.