Iris Quality in an Operational Context

James S. Doyle, Jr. and Patrick J. Flynn
Department of Computer Science and Engineering
University of Notre Dame
{jdoyle6, flynn}@nd.edu

Abstract

The accuracy of an iris biometrics system increases with the quality of the sample used for identification. All current iris biometrics systems capture data streams that must be processed to identify a single, ideal image to be used for identification. Many metrics exist to evaluate the quality of an iris image. This paper introduces a method for determining the ideal iris image from a set of iris images by using an iris-matching algorithm in a feedback loop to examine the set of true matches. This proposed method is shown to outperform other methods currently used for selecting an ideal image from a set of iris images.

Introduction

Biometrics is the use of one or more intrinsic characteristics to identify one person to the exclusion of others. There are many characteristics that can be used as biometrics, individually or combined in some manner to produce a hybrid biometric. Identification using the ear (Yan & Bowyer 2007), gait (L. Wang et al. 2003), and even body odor (Korotkaya) have been studied alongside more traditional biometrics such as face (Bowyer, Chang, & Flynn 2006), fingerprint (Chapel 1971), and iris (Bowyer, Hollingsworth, & Flynn 2008). The human iris is considered to be a mature biometric, with many real-world systems (LG 2010; Daugman & Malhas 2004; UK Border Agency 2010; Life Magazine 2010; Welte 2010) already deployed. The system designed by John Daugman (Daugman 2002) was documented to achieve a false match rate of zero percent in a particular application. In larger deployments there is still room for improvement.

Iris biometric systems can be used for identification purposes in which no claim of identity is presented, or for verification purposes where a sample and a claim of identification are supplied for the system to verify. For positive identification, both methods require that the subject be previously enrolled in the system, making him or her part of the gallery. Another sample, known as the probe, will be taken at the time of the identification attempt and compared to gallery samples. If the new probe sample closely resembles one of the gallery samples, the system reports a positive ID. If no gallery sample matches the probe sample close enough, the system will not report a match. For verification, an identity claim is also presented along with a probe sample, allowing the system to only consider samples with the requested ID. If these probe-gallery comparisons meet the threshold set for the system, then the ID claim is verified.

One method that can increase the correct recognition rate of iris biometric systems is to select the highest quality sample of a subject to be used as the representative in the gallery. Some systems have an enrollment process that takes multiple samples and then uses the highest quality sample based on certain criteria, or quality metrics, for the gallery. Samples may also be enhanced after acquisition via contrast stretching, histogram normalization, image de-blurring, or by other methods.

The Iridian LG EOU 2200 used to capture biometric samples at the University of Notre Dame is capable of outputting video as well as still images when enrolling subjects. Multiple methods have been applied to the iris videos recorded with this camera in order to extract the frame of highest quality for enrollment. For instance, the IrisBEE (Liu, Bowyer, & Flynn 2005) software, discussed in Section 4, uses sharpness to determine the focus of a frame. It computes a focus score that can be used to rank the frames, and accordingly pick the sharpest frame. It will be shown later that this approach is not optimal.

This paper suggests a new method for quality-based ranking of a set of iris biometrics samples, shows that the new method outperforms the previous quality metric, and offers possible optimizations to reduce the run-time complexity of determining the sample ranking.

Related Work

The National Institute of Standards and Technology (Tabassi, Grother, & Salamon 2010) has tentatively defined the quality of an iris image over 12 dimensions, including contrast, scale, orientation, motion blur, and sharpness. However, the definitions of these quality metrics, as well as their implementation, are still undergoing debate.

Kalka et al. (Kalka *et al.* 2006) have examined multiple quality metrics for iris images, including defocus blur, motion blur, and off-angle gaze, as well as others. Using the Dempster-Shafer theory, a fused score was calculated for all tested quality metrics. This fused score was taken as the overall quality score for an image. Using a cutoff value, it was shown that the performance of a matching system was improved when only images with scores above the cutoff

were included in the experiment.

Kalka et al. (Kalka *et al.* 2010) have extended their work to a fully-automated program to evaluate iris quality. The program creates a single, fused score, based on multiple criteria. Using this program to remove poor-quality samples from multiple datasets, they were able to demonstrate an improvement in recognition performance.

Kang and Park (Kang & Park 2005) propose a method of restoring images of poor quality due to defocus blur. The experimental method discussed in the paper showed a decrease in Hamming distances when the original image was compared to the restored image versus when the original image was compared to an artificially defocused image.

Experimental Setup

Data Acquisition

The method proposed in this paper is general, and able to be applied to any set of iris samples from the same subject. Iris videos were used in this paper as a means of capturing large amounts of same-subject samples.

All iris samples were captured using the Iridian LG EOU 2200 video camera. The LG 2200 uses three near-infrared LEDs to illuminate the eye during acquisition: one above the eye, one to the bottom left of the eye and one to the bottom right of the eye. The LG 2200 uses only one illuminant at a time to reduce spectral highlights. Software on a workstation selects one candidate frame under each lighting scenario using the proprietary IrisAccess software, of which one would be enrolled. The LG 2200 also outputs an NTSC video stream, which is recorded and encoded into a high-bit-rate MPEG-4 video. All videos captured from this device were interlaced and recorded at a constant resolution of 720x480 pixels. Figure 1 depicts the details of the iris video acquisition setup.

All videos were captured in the same windowless indoor lab under consistent lighting conditions. Subjects were supervised during acquisition to ensure proper acquisition procedure was followed. Each video was captured while the subject was being enrolled by the IrisAccess software.

Data Manipulation

The conversion of the raw data stream to video has a noticeable effect on the quality of the images. Still frames from the videos are stretched slightly along the X-axis due to digitizer timing. This was corrected by shrinking the images 5% in the X direction before experimentation began. Images recorded at 720x480 became 684x480. Additionally, a contrast-stretching algorithm was applied to all the images such that 1% of the output pixels were white and 1% were black. These two steps were helpful in improving the quality of the input data set. Figure 2 shows an example frame in original, contrast stretched, and a resized states.

Data Selection

To evaluate the performance of quality metrics for iris biometrics in an operational environment, 1000 test videos were chosen from 11751 videos acquired from 2007-2009 and stored in the BXGRID biometrics database at University of

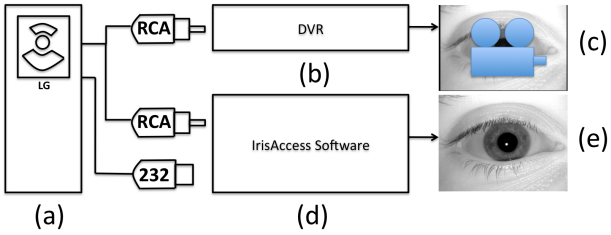

Figure 1: The LG 2200 iris camera (a) outputs one NTSC video data stream that is split and amplified using a powered splitter. The NTSC data stream is recorded by a DVR (b). This method is used to capture iris videos (c). The IrisAccess software (d) running on a workstation monitors the NTSC signal. This method is used to capture still images (e).

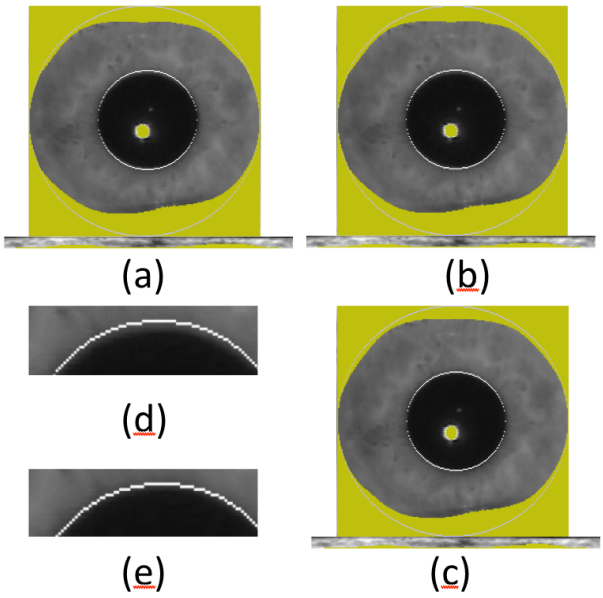

Figure 2: The three image types discussed in the paper are shown, with white circles indicating the segmentation and yellow regions representing spectral highlight removal and eyelash/eyelid masking. Image (a) is from video 05707d173 in its unmodified state. Image (b) is the contrast stretched version of (a). Image (c) is (b) scaled down by 5% in the X direction. Image (d) shows a close-up of image (a), demonstrating the segmentation error and necessity for resizing. Image (e) shows the same region after resizing.

Notre Dame (Bui *et al.* 2009). This video set included five separate videos for each of 200 subject irises. Because the time required to capture ideal frames illuminated by each light source was variable, the video lengths were not constant. The average video length was 620 frames, with a minimum of 148 frames and a maximum of 2727 frames. No effort was made to screen the data set to eliminate videos of especially poor quality, to keep the test system as close to real-world as possible.

Biometrics Software Packages

IrisBEE

IrisBEE, a modified version of the system originally developed by Masek (Masek 2003), modified by Liu (Liu, Bowyer, & Flynn 2006), released by NIST as part of the ICE challenge dataset, and further modified by Peters (Peters, Bowyer, & Flynn 2009), was used to identify iris boundaries as well as perform matching experiments. IrisBEE contains three executables, one for segmenting iris images and producing iris codes, one for performing matching experiments, and one for calculating rudimentary quality scores, based on image sharpness.

The IrisBEE IrisPreprocessor (Liu, Bowyer, & Flynn 2005) uses computer vision algorithms to detect an iris in an image. A Canny edge detector and Hough transform are used to identify the iris-pupil and iris-sclera boundaries. Active contours (Daugman 2007) are used to further refine the region borders. Two non-concentric circles are fitted to the contour to represent these two boundaries. The iris region formed by these circles is transformed into a rectangle through a polar-Cartesian conversion. Each row of the enrolled image is convolved with a one-dimensional log-Gabor filter. The complex filter response forms the iris code used for matching. A fragile bit mask (Hollingsworth, Bowyer, & Flynn 2007) is applied to allow the more stable regions of the iris code to be used in comparisons. Masking fragile bits improves the match rate, allowing for better match results than when comparing iris codes without masking fragile bits.

IrisBEE also supplies an executable for matching a gallery list of iris codes to a probe list of iris codes. The IrisMatcher outputs a fractional Hamming distance, using formula (1), for each individual comparison as well as the number of bits that were compared between the two codes, useful in normalization (Daugman 2007). All matching results from every experiment were normalized using equation (2) (Daugman 2007). Normalized fractional Hamming distances are referred to as "matching scores" throughout the rest of this paper.

For the purposes of predicting performance of a certain image in the IrisMatcher, IrisBEE provides a QualityModule executable that can rate images based on the sharpness of the whole image or just the iris region if it has been defined. The used by the QualityModule is described by Kang and Park (Kang & Park 2007). The sum of the response at each pixel of the input image is used as the image's score. The higher the score, the better that image's rank.

$$HD_{raw} = \frac{|(codeA \otimes codeB) \cap maskA \cap maskB|}{|maskA \cap maskB|} \quad (1)$$

$$HD_{norm} = 0.5 - (0.5 - HD_{raw}) \sqrt{\left(\frac{n}{900}\right)} \quad (2)$$

Neurotechnology VeriEye

A commercially-available biometrics package, Neurotechnology VeriEye (version 2.2), was also used for segmenting iris images and matching iris templates. Since VeriEye is a proprietary software package, details about the segmentation and matching algorithms are not available. The VeriEye matching algorithm reports match scores from 0 to 3235, higher scores indicating better matches. If the VeriEye matcher determines a pair of templates to be of different irises based on a threshold, it reports a match score of 0. For all experiments discussed here, this threshold was disabled to capture raw match scores, unlike the Hamming distance scores reported by IrisBEE. Input to the VeriEye matcher is order dependent, different match scores can be observed depending on the order of gallery and probe. For this paper, only one matching score was considered, with the older of the two images being the gallery image and the newer of the two being the probe image.

Smart Sensors MIRLIN

MIRLIN, another closed-source biometrics package was used to segment iris images and match iris templates, as well as to rate images based on four common quality metrics. Since MIRLIN is proprietary, specific details about its segmentation and matching algorithms as well as its quality metrics are not available. MIRLIN does provide matching scores as Hamming distances, but does not supply the number of bits used in the comparison, making normalization impossible. As a result, matching scores from MIRLIN can not be directly compared to those produced by IrisBEE. Matching scores are also symmetrical, so comparison order is not important. The four quality metrics that are discussed in this work are contrast, saturation, sharpness, and signal-to-noise ratio. MIRLIN also reports the average graylevel and occlusion percentage, but these quality metrics were not useful in classifying images since they had a very small range.

Quality Metric Experiments

Multiple quality metrics were considered: the IrisBEE QualityModule, the MIRLIN quality metrics, and the method proposed here, evaluated using IrisBEE, VeriEye, and MIRLIN. Each separate quality metric was evaluated in a similar manner so that results could be compared experimentally across the metrics.

IrisBEE Quality Module

The IrisBEE QualityModule is the current metric used at the University of Notre Dame to determine an ideal image or subset of images to be used in matching experiments from

an input set. Since the QualityModule processes individual images, the 1000 subject videos were split into individual images. The images were then segmented using the IrisPreprocessor to identify the iris region of each image. Images that failed to segment were not included in the matching experiment. After segmentation, every frame f of a single video was given a quality score f_s by the QualityModule, higher scores indicating higher quality images. The frame scores were then sorted from highest to lowest such that $f_s[i] \geq f_s[i+1] \; \forall \; i$.

To test whether this method is predictive of performance, the entire range of scores must be included in an all-vs-all matching. Due to the scale of this experiment, a subset of nine images was chosen from each video: the top-ranked image, the bottom-ranked image and seven equally spaced images in between. Selecting images in this manner controlled for the inconsistent video length of the data set. This reduced dataset was used in an all-vs-all matching experiment using the IrisMatcher. Receiver Operating Characteristic (ROC) curves were then created for each octile. These ROC curves can be found in Figure 3 and Figure 4.

With the exception of the top-ranked images, recognition performance was monotonically decreasing as QualityModule rank increased. The top-ranked images chosen by the QualityModule do not perform well in the matching experiment. As the QualityModule will recommend images with high sharpness, images with eyelid/eyelash occlusion will have artificially high scores. This causes some poor quality images to be highly ranked. Figure 5 shows a top-ranked frame with high occlusion and a lower-ranked frame more ideal for matching experiments, illustrating the drawbacks to reliance on the IrisBEE metric and other metrics that estimate image quality only.

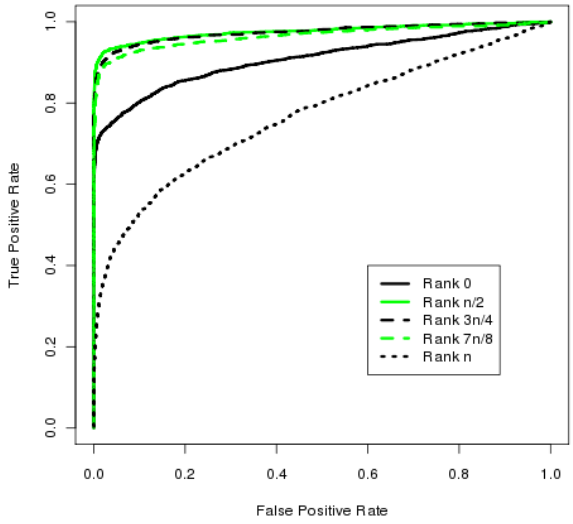

Figure 3: IrisBEE QualityModule experiment results as ROC curves, at select rank octiles. Normalized video length represented by n.

MIRLIN Quality Metrics

The 1000 subject videos were split into individual images and segmented using MIRLIN to identify the iris region of each image. Images that failed to segment were not included in the matching experiment. After segmentation, every frame f of a single video was given four quality scores $f_{contrast}$, $f_{saturation}$, $f_{sharpness}$, and f_{snr}, by MIRLIN. Four rankings were determined for each video: $f_{contrast}[i] \geq f_{contrast}[i+1] \; \forall \; i$, $f_{saturation}[i] \leq f_{saturation}[i+1] \; \forall \; i$, $f_{sharpness}[i] \geq f_{sharpness}[i+1] \; \forall \; i$, and $f_{snr}[i] \leq f_{snr}[i+1] \; \forall \; i$.

The same experimental setup as was used in the IrisBEE Quality Module experiment was used here. The same phenomenon was noticed with all four of the MIRLIN quality metrics studied. In all cases, the Rank 0 frames were outperformed by the Rank n/8 frames. ROC curves for these experiments can be found in Figured 6.

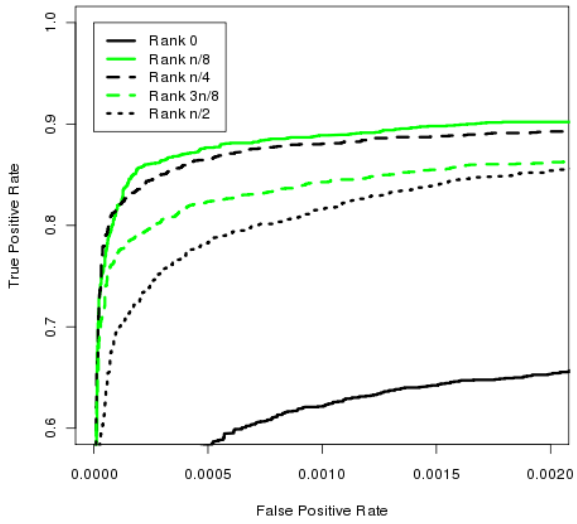

Figure 4: IrisBEE QualityModule experiment results as ROC curves, for selected octiles. Normalized video length represented by n.

IrisBEE IrisMatcher

Since the goal of this research is to find the image or set of images that performs best in matching experiments to represent a subject in a gallery, we investigated the use of the IrisMatcher itself to rate individual frames. To harness the IrisMatcher to pick an ideal representative sample, all frames of

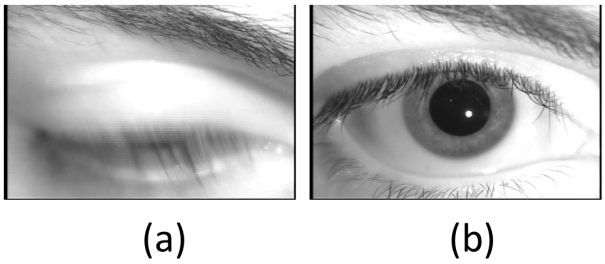

Figure 5: Sample images ranked by the QualityModule, illustrating non-intuitive quality scores from video 05697d222. Image (a) was the highest ranked image by the QualityModule. Image (b) was the 50^{th} ranked frame of 1390 images in the same video.

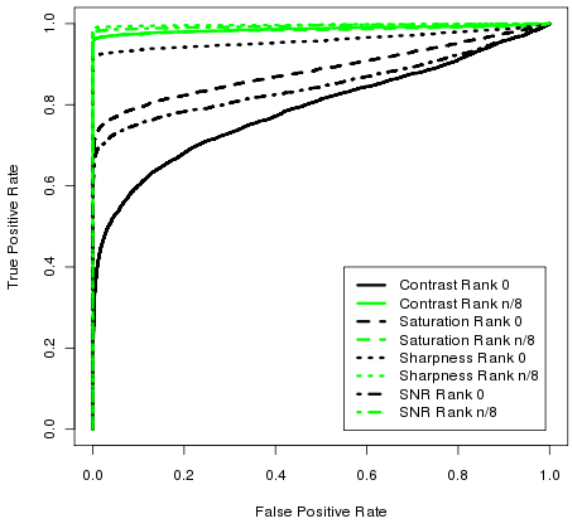

Figure 6: MIRLIN Quality Metrics experiments results as ROC curves, at select rank octiles. Normalized video length represented by n.

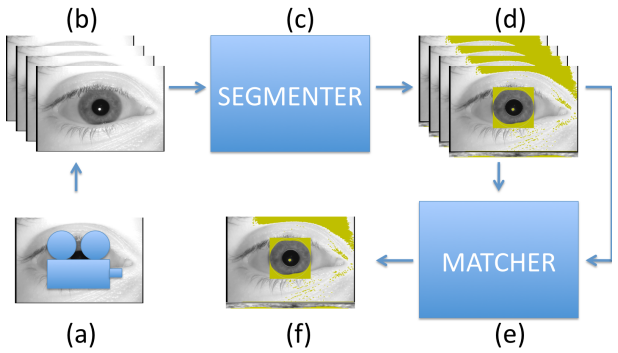

Figure 7: An iris video (a) is broken up into individual frames (b) and processed by the segmenter (c) to identify and mask out occluded regions of the iris. These iris segments (d) are then used as input into the matcher (e), which computes an all-pairs matching matrix. After some processing of the matcher results, an optimal frame can be determined (f).

the video are stored as images. The IrisBEE IrisPreprocessor was used to segment all images and to produce template files for matching. The IrisMatcher performs an all-vs-all matching of the input images for a single video, producing a fractional Hamming distance (1) for each unique pair of images. An average matching score f_s per frame f is found by averaging all matching scores resulting from comparisons involving that image. Since each iris video used in this experiment contained only one subject, all comparisons in this step were true matches. A separate matching is performed for each video. The average matching scores for each frame of a video are sorted lowest-to-highest such that $f_s[i] \leq f_s[i+1] \; \forall \; i$, since low matching scores denote more similar items or better matches. This process is illustrated in Figure 7.

Please refer to Section 5.1 for frame selection method. ROC curves can be found in Figure 8 and Figure 9.

With no exceptions, recognition performance was monotonically decreasing as IrisBEE IrisMatcher rank increased. The amount of separation between ranks in the higher-ranked half of the set was orders of magnitude smaller than the separation of the lower-ranked half.

Neurotechnology VeriEye

The IrisBEE IrisMatcher experiment was repeated using the Neurotechnology VeriEye package. All images were segmented and templates were generated using the VeriEye segmenter. The VeriEye matcher performed an all-vs-all matching of the input images for a single video, producing a matching score between 0 and 3235 for each unique pair of images.

Please refer to Section 5.1 for frame selection method. ROC curves can be found in Figure 10 and Figure 11.

With no exceptions, recognition performance was monotonically decreasing as VeriEye matcher rank increased. The amount of separation between ranks in the higher-ranked

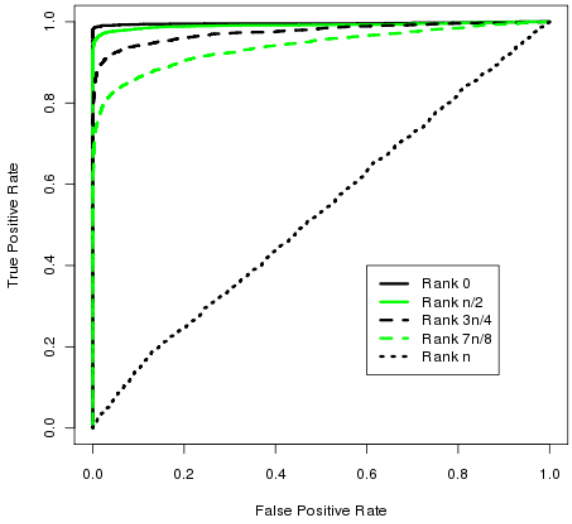

Figure 8: IrisBEE IrisMatcher experiment results as ROC curves, at selected octiles. Normalized video length represented by n.

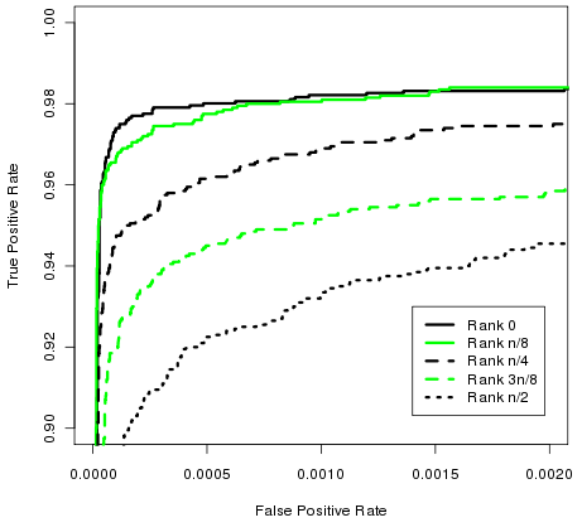

Figure 9: IrisBEE IrisMatcher experiment results as ROC curves, at selective octiles. Normalized video length represented by n.

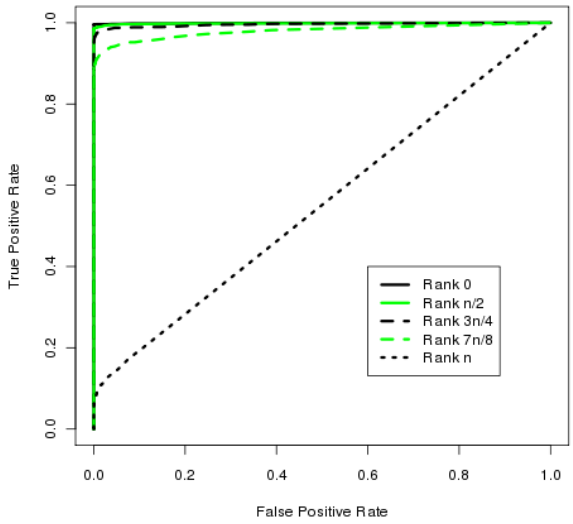

Figure 10: VeriEye matcher experiment results as ROC curves, select octiles. Normalized video length represented by n.

half of the set was orders of magnitude smaller than the separation of the lower-ranked half.

Smart Sensors MIRLIN

The experiment was again repeated using the MIRLIN package. All images were segmented and templates were generated using the MIRLIN $-get$ command. The MIRLIN matcher performed an all-vs-all matching of the input images for a single video, using the MIRLIN $-compare$ command, producing a matching score in the range $\{0,1\}$ for each unique pair of images.

Please refer to section 5.1 for frame selection method. ROC curves can be found in Figure 12 and Figure 13.

As was the case with IrisBEE and VeriEye, with no exceptions, recognition performance was monotonically decreasing as MIRLIN matcher rank increased. The amount of separation between ranks in the higher-ranked half of the set was orders of magnitude smaller than the separation of the lower-ranked half.

Quality Metric Comparison

For all experiments, there is an ordering, with higher ranked frames performing better in all cases except for the top frame reported by the IrisBEE and MIRLIN quality metrics. Poor performance of the top-ranked frame can be explained by the mechanism in which these quality metrics rank images. Since the quality metrics use image analysis techniques on the iris texture to rate a frame it can be heavily influenced by eyelashes or spectral highlights that were not properly masked, or other artifacts present in an image. These artifacts in turn produce noisy samples as they are blocking

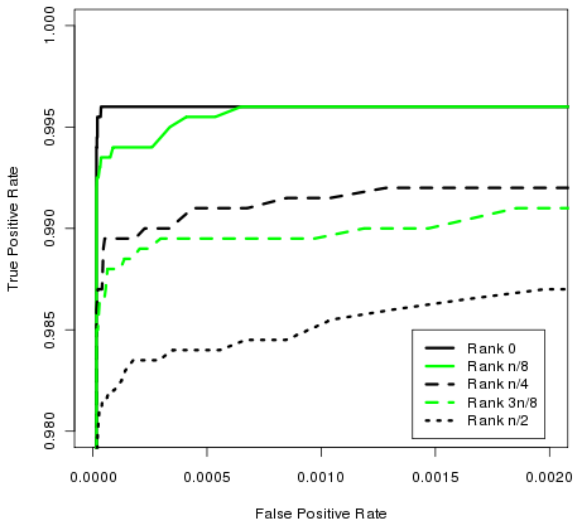

Figure 11: VeriEye matcher experiment results as ROC curves, top octiles. Normalized video length represented by n.

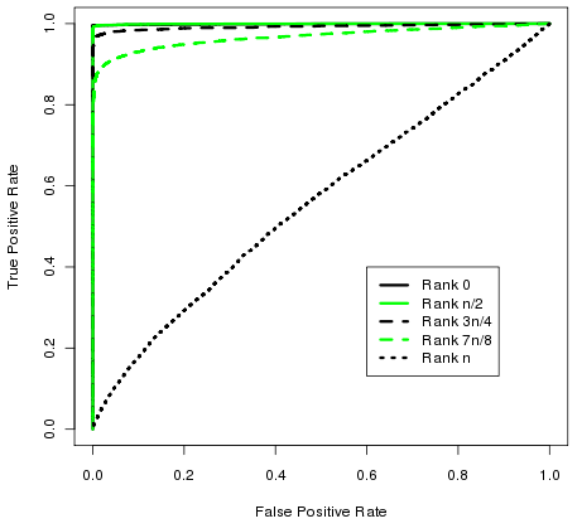

Figure 12: MIRLIN matcher experiment results as ROC curves, select octiles. Normalized video length represented by n.

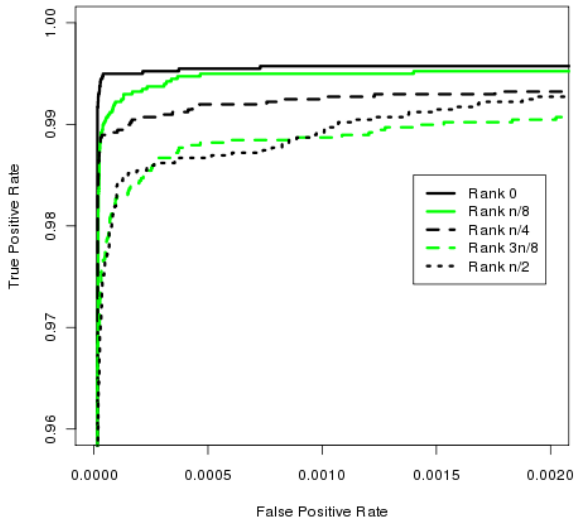

Figure 13: MIRLIN matcher experiment results as ROC curves, top octiles. Normalized video length represented by n.

parts of the iris texture from being compared, artificially skewing the match score higher than it should be for a known match (Bowyer, Hollingsworth, & Flynn 2008). However, even the best images from these quality metrics did not perform as well as the self-matching experiments. The self-matching image selection process can eliminate samples with these artifacts from being used in biometric comparisons by minimizing (or maximizing in the case of VeriEye) the average match scores.

The ordering seen in the ROC curves indicates that the intra-video IrisBEE, VeriEye, and MIRLIN match scores are predictive of inter-video matching performance. Figure 14 shows the ROC curve for the top-ranked frame from each of the metrics, as well as the n/8-ranked frames from each of the quality metrics, as these were the highest performing ranks for the quality metrics.

Application

Data from the LG 2200 iris camera was used in this paper because it allows video streams to be captured easily. Other LG iris cameras, as well as iris cameras from other manufacturers, do not allow video information to be output from the device. However, the use of video in this paper was merely for convenience. The self-matching algorithm could be applied to any set of data captured by the same sensor, including a small set of still images captured by a newer iris camera.

Although production versions of most iris cameras do not allow video to be captured from the device, the proprietary software that interfaces with the camera does capture video information. This method could be applied to the data steam that is processed by the proprietary software. However, as

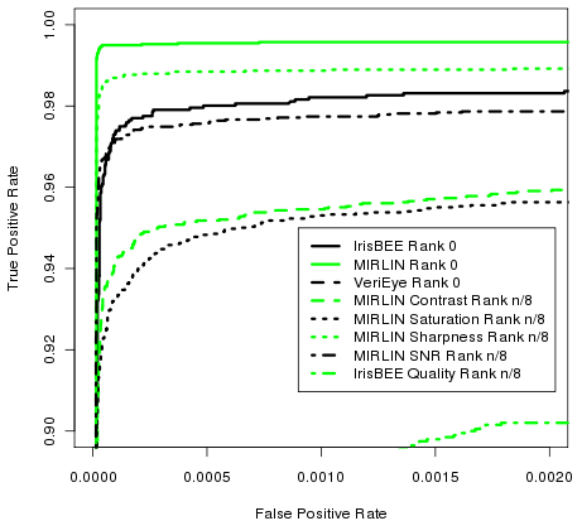

Figure 14: ROC curves of top octiles from all metrics, plus QualityModule and MIRLIN Quality octile n/8 which displayed best performance.

this method is somewhat time consuming, it may only be feasible to apply it during the enrollment phase. Performing this analysis for every probe would delay the response from the system by an unacceptably large amount.

Conclusions

It has been shown through empirical analysis that this method selects an ideal representative sample from a set of same-subject samples. This method outperforms a sharpness-based metric used currently and can be used with a commercially available system.

References

[Bowyer, Chang, & Flynn 2006] Bowyer, K. W.; Chang, K.; and Flynn, P. 2006. A survey of approaches and challenges in 3d and multi-modal 3d + 2d face recognition. *Computer Vision and Image Understanding* 101(1):1 – 15.

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

[Bui *et al.* 2009] Bui, H.; Kelly, M.; Lyon, C.; Pasquier, M.; Thomas, D.; Flynn, P.; and Thain, D. 2009. Experience with BXGrid: a data repository and computing grid for biometrics research. *Cluster Computing* 12:373–386. 10.1007/s10586-009-0098-7.

[Chapel 1971] Chapel, C. 1971. *Fingerprinting: A Manual of Identification*. Coward McCann.

[Daugman & Malhas 2004] Daugman, J., and Malhas, I. 2004. Iris recognition border-crossing system in the UAE. http://www.cl.cam.ac.uk/~jgd1000/UAEdeployment.pdf.

[Daugman 2002] Daugman, J. 2002. How iris recognition works. volume 1, I–33 – I–36 vol.1.

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

[Hollingsworth, Bowyer, & Flynn 2007] Hollingsworth, K.; Bowyer, K.; and Flynn, P. 2007. All iris code bits are not created equal. 1 –6.

[Kalka *et al.* 2006] Kalka, N. D.; Zuo, J.; Schmid, N. A.; and Cukic, B. 2006. Image quality assessment for iris biometric. volume 6202, 62020D. SPIE.

[Kalka *et al.* 2010] Kalka, N.; Zuo, J.; Schmid, N.; and Cukic, B. 2010. Estimating and fusing quality factors for iris biometric images. *Systems, Man and Cybernetics, Part A: Systems and Humans, IEEE Transactions on* 40(3):509 –524.

[Kang & Park 2005] Kang, B., and Park, K. 2005. A study on iris image restoration. In Kanade, T.; Jain, A.; and Ratha, N., eds., *Audio- and Video-Based Biometric Person Authentication*, volume 3546 of *Lecture Notes in Computer Science*. Springer Berlin / Heidelberg. 31–40. 10.1007/115279234.

[Kang & Park 2007] Kang, B. J., and Park, K. R. 2007. Real-time image restoration for iris recognition systems. *Systems, Man, and Cybernetics, Part B: Cybernetics, IEEE Transactions on* 37(6):1555 –1566.

[Korotkaya] Korotkaya, Z. Biometric person authentication: Odor. Lappeenranta University of Technology.

[L. Wang et al. 2003] L. Wang et al. 2003. Silhouette analysis-based gait recognition for human identification. *IIEEE Trans. Pattern Analysis and Machine Intelligence* 25(1).

[LG 2010] LG. 2010. IrisID in Action. http://irisid.com/ps/inaction/index.htm.

[Life Magazine 2010] Life Magazine. 2010. Iris identification used by parents and teachers to protect children in nj elementary school. http://life.com/image/1965902.

[Liu, Bowyer, & Flynn 2005] Liu, X.; Bowyer, K.; and Flynn, P. 2005. Experiments with an improved iris segmentation algorithm. 118 – 123.

[Liu, Bowyer, & Flynn 2006] Liu, X.; Bowyer, K. W.; and Flynn, P. J. 2006. *Optimizations in Iris Recognition*. Ph.D. Dissertation, University of Notre Dame.

[Masek 2003] Masek, L. 2003. Recognition of human iris patterns for biometric identification. Technical report, The University of Western Australia.

[Peters, Bowyer, & Flynn 2009] Peters, T.; Bowyer, K. W.; and Flynn, P. J. 2009. Effects of segmentation routine and acquisition environment on iris recognition. Master's thesis, University of Notre Dame.

[Tabassi, Grother, & Salamon 2010] Tabassi, E.; Grother, P.; and Salamon, W. 2010. Iris quality calibration and evaluation 2010. *IREX II IQCE*.

[UK Border Agency 2010] UK Border Agency. 2010. Using the Iris Recognition Immigration System (IRIS). http://www.ukba.homeoffice.gov.uk/travellingtotheuk/Enteringtheuk/usingiris.

[Welte 2010] Welte, M. S. 2010. Prison system looks to iris biometrics for inmate release. http://www.securityinfowatch.com/Government+%2526+Public+Buildings/1314995.

[Yan & Bowyer 2007] Yan, P., and Bowyer, K. W. 2007. Biometric recognition using 3d ear shape. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 29:1297–1308.