Fusion of Face and Iris Biometrics from a Stand-Off Video Sensor

Ryan Connaughton and **Kevin W. Bowyer** and **Patrick Flynn**
University of Notre Dame
Department of Computer Science and Engineering
Notre Dame, IN 46556

Abstract

Multi-biometrics, or the fusion of more than one biometric modality, sample, sensor, or algorithm, is quickly gaining popularity as a method of improving biometric system performance and robustness. Despite the recent growth in multi-biometrics research, little investigation has been done to explore the possibility of achieving multi-modal fusion from a single sensor. This approach to multi-biometrics has numerous advantages, including the potential for increased recognition rates, while still minimizing sensor cost and acquisition times. In this work, experiments are presented which successfully combine multiple samples of face and iris biometrics obtained from a single stand-off and on-the-move video sensor. Several fusion techniques are explored, with the best recognition rates achieved by using a weighted summation of face and iris match scores. The fusion results out-perform either single-modality approach, and the proposed multi-biometric framework represents a viable and natural extension to the stand-off iris sensor used to acquire subject data.

Introduction

The practice of using more than one biometric modality, sample, sensor, or algorithm to achieve recognition, commonly referred to as multi-biometrics, is a technique that is rapidly gaining popularity. By incorporating multi-biometrics into the recognition process, many of the shortcomings of traditional single-biometric systems can be alleviated and overall recognition accuracy can be improved. Multi-biometrics can inherently increase system robustness by removing the dependency on one particular biometric approach. Further, a system that utilizes more than one biometric feature or matcher may be more difficult to deliberately spoof (Ross, Nandakumar, & Jain 2006). Systems that make use of multiple biometric features can also provide redundancy that may lower failure-to-acquire rates.

While research into multi-biometrics has received a large increase in attention over recent years, the task of fusing multiple biometric modalities from a single sensor remains an under-studied challenge. Due to a lack of available multi-modal data, many current experiments in multi-biometrics create "chimeric" datasets, in which samples of one biometric modality from one set of subjects are arbitrarily paired with a second biometric modality from a separate set of subjects in order to simulate a multi-biometric scenario (Bowyer *et al.* 2006). This approach, though useful for preliminary experimentation, may mask unknown dependencies between modalities. Further, chimeric datasets simulate a multi-biometric scenario in which samples of each modality are acquired independently. In practice, it is much more desirable to simultaneously acquire multiple modalities from a single sensor if possible for cost and usability reasons. This work presents a multi-biometric system which simultaneously acquires face and iris information under near-infrared (NIR) illumination using the Iris on the Move (IOM) sensor, which is composed of an array of three identical video cameras with timed NIR illumination (Matey *et al.* 2006). The face and iris information for each subject is combined to improve recognition rates beyond the observed recognition rates for either isolated biometric.

Background and Related Work

There are four general approaches to multi-biometric system design. In the *multi-sample* approach, multiple samples (e.g. images) of the same biometric modality are acquired and processed. In the *multi-sensor* approach, the same modality is sampled several times, using different sensors for each acquisition. In the *multi-algorithm* approach, each biometric sample is matched using multiple matching algorithms and the results are fused. Finally, *multi-modal* systems acquire samples of more than one biometric trait (e.g. iris and face) for matching. Additionally, some systems represent a hybrid of these approaches by adding redundancy at multiple stages of the recognition process.

There are several levels at which fusion can occur in a multi-biometric system. Using *signal-level* fusion, multiple samples may be combined together to create one superior sample (as in super-resolution techniques). Alternatively, features can be extracted from each biometric sample, and *feature-level* fusion can be used to condense all of the features into a single biometric signature. With *score-level* fusion, each sample is processed and matched separately, and the resulting match scores for each sample are combined into one final match score. *Rank-level* fusion combines match rankings, rather than the scores for each sample, into a final ranking to determine the best match. Finally, *decision-level* fusion applies a matcher to each biometric sample to determine whether or not each comparison is a match, and the response of each matcher is fused using Boolean operators, a voting scheme, or some similar method.

The fusion of face and iris modalities is a biometric approach that has gained increasing attention over the past decade, likely due to the popularity of the individual modalities and the natural connection between them. Despite this recent trend, very few studies have been done on fusion of face and iris biometrics from a single sensor.

The most common method of multi-biometric fusion is score-level fusion. Zhang et al. approach the problem of fusing face and iris biometrics under near-infrared lighting using a single sensor (Zhang *et al.* 2007). Frontal face images are acquired using a 10 megapixel CCD camera. Eye detection and face alignment are performed using Local Bit Pattern histogram matching as described in Li et al. (Li *et al.* 2006). The eigenface algorithm and Daugman's algorithm are used to perform face and iris recognition, respectively, and score-level fusion is accomplished via the sum and product rules after min-max normalization. Numerous other score-level fusion approaches have been tested on chimeric datasets. Chen and Te Chu use an unweighted average of the outputs of matchers based on neural networks (Chen & Te Chu 2005). Wang et al. test weighted average, linear discriminant analysis, and neural networks for score fusion (Wang, Tan, & Jain 2003).

Another common approach to biometric fusion is feature-level fusion through concatenation. Rattani and Tistarelli compute SIFT features for chimeric face and iris images and concatenate the resulting feature vectors (Rattani & Tistarelli 2009). The number of matching SIFT features between two vectors (measured by Euclidean distance) is used as a match score for that comparison. Son and Lee extract features for face and iris images based on a Daubechies wavelet transform (Son & Lee). Concatenation is used to form a joint feature vector, and Euclidean distance between feature vectors is used to generate match scores.

Approach

To facilitate the fusion of face and iris biometrics from a single sensor, we selected the Iris on the Move (IOM) sensor for data acquisition. The IOM is a sensor designed for high-throughput stand-off iris recognition (Matey *et al.* 2006). The IOM features a portal which subjects walk through at normal walking pace. As a subject passes through the portal, the subject is illuminated with near-infrared (NIR) LED's, and frontal video is captured by an array of three vertically-arranged, fixed-focus cameras equiped with NIR filters. The presence of multiple cameras allows the system to handle a larger range of subject heights. Though the sensor is intended for iris image acquisition, the face is typically captured as well. While the sides of the portal help to direct subjects into the field of view of the cameras, it is possible for subjects to stray partially out of the video frames, leading to frames with partial faces or only one visible iris. Figure 1 shows corresponding frames from each of the three IOM cameras while a subject passes through the in-focus region of the IOM. Each frame captured by one of the IOM cameras is a 2048 by 2048 pixel grayscale image. A typical iris acquired by the system is approximately 120 pixels in diameter.

Figure 1: Example of corresponding frames from the IOM as the subject passes through the in-focus region of the portal. The left image shows a frame from the top camera, the middle image shows a frame from the middle camera, and the right shows a frame from the bottom camera.

The general steps used in this work to combine face and iris biometrics from the IOM sensor are outlined in Figure 2. As previously described, when a subject passes through the IOM portal, three videos are collected, with one video coming from each of the IOM cameras. In a preprocessing step, the corresponding frames of the three videos are stitched together to create one single video. Next, a series of detection phases are used to locate whole faces and eyes in each frame. Matching is then performed on each face and iris independently, and the results are fused using several different techniques.

Preprocessing

In order to increase the likelihood of a whole face being captured for each subject, the three videos from each IOM acquisition are "stitched" together to combine corresponding frames. As can be seen in Figure 1, there is significant vertical overlap between the top and middle cameras, as well as between the middle and bottom cameras. Due to imperfect calibration of the individual cameras, some horizontal misalignment between the cameras is also present.

A template-matching approach is taken to determine the desired translation to align frames from adjacent cameras. Specifically, the bottom portion of the top frame is cropped and used as a template. This template is then matched against the upper half of the middle frame, and the best match is selected as the desired alignment. This process is repeated for the bottom camera, where the template is created from the top portion of the bottom frame and matched against the lower half of the middle frame.

Finally, noticeable illumination differences were observed between corresponding frames from different cameras, likely due to mis-calibration. To account for this discrepancy, histogram matching is used to match the top and bottom frame to the illumination observed in the middle frame. Figure 3 shows the intermediate and final results of the stitching procedure for an example frame.

Face Detection

Once the frame stitching is completed, the next step in the preprocessing phase is to detect a face in each frame. To accomplish this task, the OpenCV implementation of the Viola-Jones cascade face detector is used (Bradski &

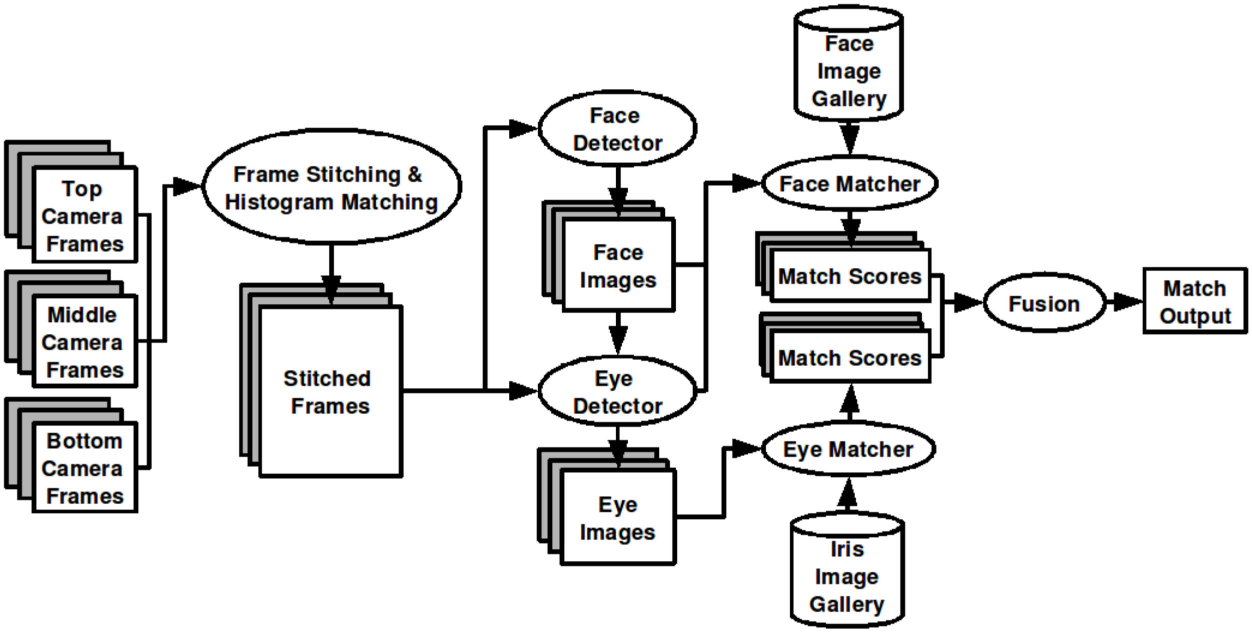

Figure 2: A diagram of the pipeline used in the proposed multi-biometric system.

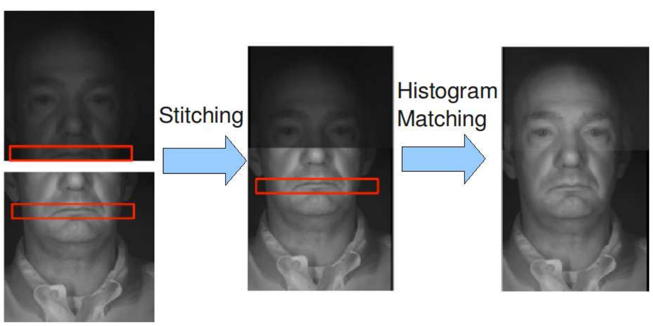

Figure 3: An example of the progression during alignment between corresponding frames from the top and middle camera. The top left image is the frame from the top camera with the template marked in red. The bottom left image is the frame from the middle camera, with the matched region marked in red. The middle image is the composite image, with the frame from the top camera cropped and padded. The overlapping region is indicated. The right image shows the final stitching results after histogram matching. A similar approach is used to stitch the frame from the bottom camera to the middle frame.

Kaehler 2008), (Viola & Jones 2001). The detector was trained on whole faces, and thus may or may not detect faces which lie only partially within the field of view of the camera.

Eye Detection

The purpose of the eye detection phase is twofold. The primary goal is to detect any eyes present in each frame for iris matching. However, the locations of the eyes that are detected in the faces produced by the face detector are also used for an alignment phase during face matching. A template matching approach is adopted for eye detection. The template used to search for eyes in each frame is based on the specular highlights generated by the reflection of the IOM LEDs.

The eye detection is completed in two phases. First, the template matching is performed on the upper left and upper right quadrants of each face detected by the face detector. This approach guarantees that each detected face will have two eye locations estimated as well.

Because it is possible for eyes to be detected in frames where whole faces were not present (or in frames where the face detector failed to detect the face), a second round of template matching is performed on any stitched frame where a face was not detected. In these frames, the location of the partial face can be crudely estimated by examining the sums of the rows and columns of the image. Once the partial face region has been estimated, the template matching is performed twice to identify the two best eye locations. Finally, the detected eyes are cropped from the corresponding location in the *original* frames to remove any possible artifacts caused by the histogram matching in the stitching phase. In cases where the detected eye is located in the overlapping region between two cameras, the eye is cropped from *both* camera frames.

Face Matching

In this work, Colorado State University's implementation of the eigenface algorithm is used for face matching (Colorado State University 2010), (Turk & Pentland 1991). To achieve alignment with the training set, the probe face images are normalized using the eye centers detected by the eye detector. The Mahalanobis cosine metric is used to compute the distance between two feature vectors. Using this metric, match scores can range from -1.0 to 1.0, with -1.0 being a perfect score. The output of the face matcher stage of the pipeline is a distance for every comparison between each probe face image and gallery face image.

Iris Matching

For the iris matcher, a modified version of Daugman's algorithm is used to compare each probe iris image to the gallery (Daugman 2002). The normalized fractional Hamming distance, referred to simply as the Hamming distance in the rest of this work, ranges from 0.0 to 1.0, with 0.0 being a perfect match. The Hamming distance is normalized to adjust low Hamming distances that occur for comparisons that used relatively few bits. The output of the iris matcher stage of the pipeline is a Hamming distance for every comparison between each probe eye image and gallery iris image.

Fusion

In this framework, there is both a multi-sample (i.e. several faces from each video) and a multi-modal (i.e. both iris and face samples from each video) dimension to problem. Consequently, there are many methods which could be used to combine the face and iris biometrics from each video. Several fusion techniques are considered at both the score and rank-level.

The first method considers only one biometric modality in the fusion process, and makes use only of the multi-sample dimension of the problem by taking the minimum score for a given modality. For example, in the MinIris approach, the minimum score for all of the iris comparisons from a given video is reported as the best match. Similarly, the MinFace approach takes the minimum score for all of the face comparisons from a given video to determine the best match. Equations 1 and 2 express the MinIris and MinFace fusion rules, respectively, for a given probe video,

$$MinIris = Min\{I_{i,j} | i = 1...n, j = 1...G\} \tag{1}$$

$$MinFace = Min\{F_{i,j} | i = 1...m, j = 1...G\} \tag{2}$$

where n and m are the number of irises and faces detected in the video, respectively, G is the number of gallery subjects, $I_{i,j}$ is the Hamming distance between the i-th iris and the j-th gallery subject, and $F_{i,j}$ is the score for the comparison between the i-th face and the j-th gallery subject.

The next type of fusion method considered is rank-level fusion, and can incorporate face, iris, or both modalities into the decision process. A Borda count is used to determine a best match across the desired biometric modalities. In a Borda count, the scores for all comparisons from a given sample are sorted such that the first rank corresponds to the best score for that sample. Each sample then casts votes for the top v ranked subjects, where the weight of each vote is inversely proportionate to rank number. Each sample votes in this manner, and the gallery subject with the most votes is taken to be the best match. In these experiments, the BordaIris method considers only the iris scores to perform fusion, and the BordaFace method considers only face scores. The BordaBoth method allows both face and iris samples to vote, with v votes being cast by each iris and face sample.

Two vote weighting schemes are tested for the BordaIris, BordaFace, and BordaBoth fusion methods. In the Linear approach, the vote weight is linearly proportional to the rank; specifically, the weight associated with the rank-n match is described by the equation

$$VoteWeight_n = v + 2 - n \tag{3}$$

and v represents the total number of votes cast by each biometric sample. In the Exponential approach, the weight of the vote is exponentially related to the rank. Specifically, the weight associated with the rank-n match is described by the equation

$$VoteWeight_n = 2^{v-n} \tag{4}$$

The third fusion method again uses score-level fusion, implementing a weighted summation of the iris and face scores. The summation rule can be expressed as Equation 5 for a given probe video,

$$SumScore_k = \frac{\alpha * \sum_{i=1}^{n} FNorm_{i,k} + \beta * \sum_{j=1}^{m} (1 - I_{j,k})}{\alpha * n + \beta * m} \tag{5}$$

where n and m are the number of irises and faces detected in the video, respectively, $I_{j,k}$ is the Hamming distance between the j-th iris and the k-th gallery subject, and $FNorm_{i,k}$ is the normalized score for the comparison between the i-th face and the k-th gallery subject. Each face score $F_{i,k}$ is normalized according to the expression

$$FNorm_{i,k} = 1 - \frac{F_{i,k} + 1}{2} \tag{6}$$

so that $0 \leq FNorm_{i,k} \leq 1$ and 1 is a perfect match. In Equation 5, α and β are coefficients used to weight the face and iris biometrics, respectively. In the presented work, $\alpha = 1 - \beta$ for simplicity. In Equation 5, $SumScore_k$ represents the final match score for the given probe video with gallery subject k; the best match score can be determined by finding the maximum $SumScore_k$ for all k. SumIris is the special case where $\alpha = 0$ and $\beta = 1$, which corresponds to summing only the iris scores to determine the best match. Similarly, SumFace is the case where $\alpha = 1$ and $\beta = 0$, and equates to summing only the normalized face scores.

Table 1: DETAILED DETECTION RESULTS

Modalities Detected	Frame Count	Video Count
Left Iris	1,447 (5.1%)	35 (1.9%)
Right Iris	2,104 (7.4%)	46 (2.4%)
Face	900 (3.2%)	2 (0.1%)
Left & Right Irises	2,495 (8.8%)	209 (11.1%)
Face & Left Iris	1,411 (5.0%)	34 (1.8%)
Face & Right Iris	724 (2.6%)	27 (1.4%)
Face & Both Irises	6,798 (24.0%)	1,522 (80.6%)
None	12,502 (44.1%)	11 (0.6%)

Experiments

Dataset

The multi-biometric system being presented was tested on a probe dataset of 1,886 IOM video sets. Note that here a video "set" refers to the corresponding videos from each of the three IOM cameras, so the dataset is comprised of 5,658 videos in total. The 1,886 videos spanned 363 unique subjects, with an average of about five videos per subject. The most frequently occurring probe subject had 15 videos in the probe set, and the least frequently occurring had one probe video.

The iris gallery contained one left eye and one right eye for each of the 363 gallery subjects. The gallery images were acquired using the LG IrisAccess 4000 (LG Iris 2010), a high-quality iris acquisition camera, and the gallery was manually screened for good quality and segmentation.

The face gallery contained one full face image for each of the 363 subjects. The gallery images were acquired using the IOM. Each of the 363 subjects in the study had an additional IOM video set acquired in which the presence of a whole face was verified manually. The frames were stitched using the process previously described, and then the best frame was manually selected and the coordinates of the eye centers were manually annotated for alignment. The PCA training was performed on the face image gallery.

Detection Results

Across the entire dataset, 14,829 left irises and 14,711 right irises were detected and successfully segmented, and 9,833 faces were detected with valid eye locations for alignment. In this context, "successful segmentation" simply means that the iris segmentation routine returned pupil and limbic boundaries; it does *not* guarantee correctness. On average, 15.7 ($\sigma = 8.1$) irises, 5.2 ($\sigma = 3.7$) faces, and 20.9 ($\sigma = 20.9$) of either biometric samples were found in each video.

Table 1 provides a breakdown of the detection results by frame and video. The 1,886 videos were composed of a total of 28,381 frames. From Table 1 it can be seen that while a large number of frames (44.1%) contained no detected features, a much larger percentage of the probe *videos* (99.3%) had at least one biometric feature detected. Further, we see that the majority (80.6%) of the probe videos contained samples of face and both iris features.

Matching Results

Figure 4 shows the match and non-match score distributions for all 9,833 detected faces. The mean match score was -0.281 with a standard deviation of 0.213, while the mean non-match score was 0.000 with a standard deviation of 0.676. If each face were treated independently, the rank-one recognition achieved for the 9,833 probes faces would be 51.6% (5,073/9,833) recognition.

The results from the left and right irises were aggregated, and Figure 5 shows the match and non-match score distributions. The mean match score was 0.398 with a standard deviation of 0.053, while the mean non-match score was 0.449 with a standard deviation of 0.013. Figure 5 shows a significant number of match comparisons with fairly high scores. Upon examination of the data, it was found that most of these scores arise from incorrect segmentation. In some cases, these high match scores were caused by severe image defocus. Additionally, there are some false positives from the eye detector (non-eye regions) that contain features that resemble pupil and limbic boundaries according to the segmentation routine. If each iris image were treated independently, the rank-one recognition achieved for all of the probe irises would be 46.6% (13,556/29,112) recognition.

Fusion Results

The results of the iris and face matchers were combined using each of the methods previously described. The rank-one recognition rates achieved by each fusion approach are shown in Table 2. In the fusion methods based on Borda-counts, the number of votes given to each sample was varied between 1 and 363 (though all samples were given the same number of votes for any given fusion experiment), and the best results for each approach are presented. Similarly, results from the optimal tested values of α and β are presented.

Summarizing, the best single-modality fusion approach was the SumIris approach, which achieved an 87.8% rank-one recognition rate. The SumBoth approach achieved the overall highest recognition rate (93.2%), and all multi-modal fusion approaches achieved higher recognition rates than the fusion methods based on a single modality.

Figure 6 shows the ROC curves for the best SumBoth and BordaBoth approaches, as well as the MinIris, MinFace, SumFace, and SumIris results for comparison. From this graph, it is clear the the BordaBoth and SumBoth approaches out-perform the single-modality fusion methods. Interestingly, while SumBoth achieved the highest rank-one recognition rate, Figure 6 shows that the BordaBoth fusion technique performs better at false positive rates less than 0.06.

In general, the videos that failed to match correctly typically had relatively few face and iris features detected. While the iris proved to be the more accurate of the two modalities in the multi-sample fusion scenarios, Figure 5 indicates that many of the iris features detected are of poor quality, represent false detections from the eye detector, or failed to segment correctly. While the fusion techniques in these experiments were able to overcome these challenges when enough samples were present, videos in which a small

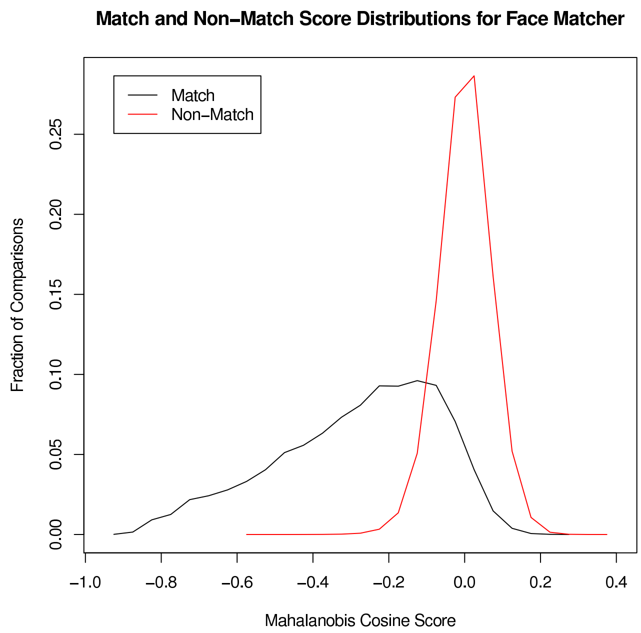

Figure 4: The match and non-match score distributions for the face features from the entire probe dataset.

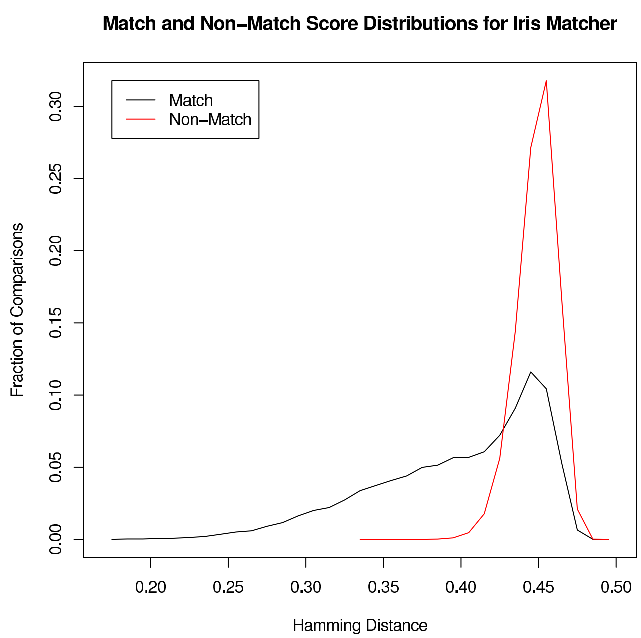

Figure 5: The match and non-match score distributions for the left and right iris features from the entire probe dataset.

Table 2: RANK ONE RECOGNITION RATES FOR FUSION APPROACHES

Approach	Fusion Parameters	Rank-One (Raw)
MinIris		86.7% (1,635/1,886)
MinFace		62.6% (1,180/1,886)
BordaIris-Linear	$v = 3$	86.4% (1,629/1,886)
BordaIris-Exponential	$v = 20$	86.8% (1,637/1,886)
BordaFace-Linear	$v = 3$	58.9% (1,110/1,886)
BordaFace-Exponential	$v = 5$	59.3% (1,118/1,886)
BordaBoth-Linear	$v = 10$	91.7% (1,729/1,886)
BordaBoth-Exponential	$v = 10$	92.0% (1,735/1,886)
SumIris	$\alpha = 0.0, \beta = 1.0$	87.8% (1,656/1,886)
SumFace	$\alpha = 1.0, \beta = 0.0$	61.3% (1,156/1,886)
SumBoth	$\alpha = 0.3, \beta = 0.7$	93.2% (1,757/1,886)

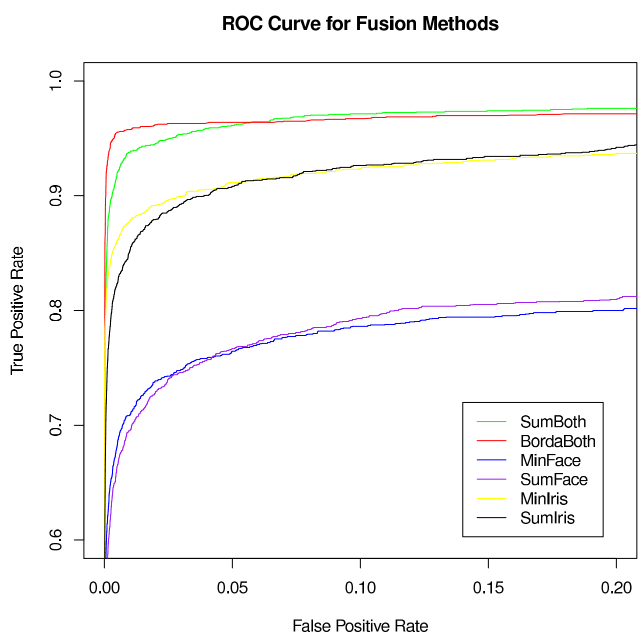

Figure 6: ROC curves for the various fusion methods using the optimal tested parameters for each. The BordaBoth method shown is the BordaBoth-Exponential method.

number of faces and iris are detected are much less likely to be correctly matched.

Conclusions and Future Work

This work presents an investigation into the fusion of face and iris biometrics from a single sensor, a surprisingly understudied problem in current literature. We present a multi-biometrics framework that utilizes both multi-sample and multi-modal fusion techniques to improve recognition rates from a single sensor. The multi-biometric system is tested on a non-chimeric dataset of over 1,886 videos spanning 363 subjects. This represents one of the largest genuine multi-modal experiments that has been conducted to date. Face and iris biometric samples extracted from videos produced from the Iris on the Move sensor were combined using several different fusion methods. In these experiments, the combination of face and iris biometrics via match score summation yielded a 5.4% increase in recognition rate over the best single-modality approach that was tested, while a modified Borda count approach performed best at lower false positive rates (< 0.06).

The multi-biometrics system we propose exploits the face information collected by the IOM, a sensor that is intended for iris recognition purposes, with no modifications to the sensor and no increase in probe data acquisition time. The resulting system is less likely to experience failures to acquire, and the use of multiple modalities could allow the system to identify subjects with incomplete gallery data. This approach could be extended to operate on other stand-off iris sensors, which often detect the face as a preliminary step to iris image acquisition.

In the future, new methods of fusion and matching will be explored, quality metrics and partial-face matching will be introduced, and run-time analysis will be conducted.

Acknowledgments

Datasets used in this work were acquired under funding from the National Science Foundation under grant CNS01-30839, by the Central Intelligence Agency, and by the Technical Support Working Group under US Army Contract W91CRB-08-C-0093. The authors are currently supported by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA), through the Army Research Laboratory (ARL). The views and conclusions contained in this document are those of the authors and should not be interpreted as representing official policies, either expressed or implied, of IARPA, the ODNI, the Army Research Laboratory, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation herein.

References

Bowyer, K. W.; Chang, K. I.; Yan, P.; Flynn, P. J.; Hansley, E.; and Sarkar, S. 2006. Multi-modal biometrics: An overview. Presented at the Second Workshop on Multi-Modal User Authentication (MMUA 2006).

Bradski, G., and Kaehler, A. 2008. *Learning OpenCV*. O'Reilly Media, Inc.

Chen, C.-H., and Te Chu, C. 2005. Fusion of face and iris features for multimodal biometrics. In Zhang, D., and Jain, A., eds., *Advances in Biometrics*, volume 3832 of *Lecture Notes in Computer Science*. Springer Berlin / Heidelberg. 571–580.

Colorado State University. 2010. Evaluation of face recognition algorithms. http://www.cs.colostate.edu/evalfacerec/algorithms5.html.

Daugman, J. 2002. How iris recognition works. In *2002 International Conference on Image Processing*, volume 1, 33–36.

LG Iris. 2010. LG Iris products and solutions. http://www.lgiris.com/ps/products/irisaccess4000.htm.

Li, S. Z.; Chu, R.; Ao, M.; Zhang, L.; and He, R. 2006. Highly accurate and fast face recognition using near infrared images. In *International Conference on Biometrics (ICB 2006)*, 151–158.

Matey, J.; Naroditsky, O.; Hanna, K.; Kolczynski, R.; LoIacono, D.; Mangru, S.; Tinker, M.; Zappia, T.; and Zhao, W. 2006. Iris on the move: Acquisition of images for iris recognition in less constrained environments. In *Proceedings of the IEEE*, volume 94. 1936–1947.

Rattani, A., and Tistarelli, M. 2009. Robust multi-modal and multi-unit feature level fusion of face and iris biometrics. In Tistarelli, M., and Nixon, M., eds., *Advances in Biometrics*, volume 5558 of *Lecture Notes in Computer Science*. Springer Berlin / Heidelberg. 960–969.

Ross, A. A.; Nandakumar, K.; and Jain, A. K. 2006. *Handbook of Multibiometrics*. Springer Science and Business Media.

Son, B., and Lee, Y. Biometric authentication system using reduced joint feature vector of iris and face. In Kanade, T.; Jain, A.; and Ratha, N., eds., *6th International Conference on Audio- and Video-Based Biometric Person Authentication (AVBPA'03)*, Lecture Notes in Computer Science. Springer Berlin / Heidelberg.

Turk, M., and Pentland, A. 1991. Face recognition using eigenfaces. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR '91)*, 586–591.

Viola, P., and Jones, M. 2001. Rapid object detection using a boosted cascade of simple features. In *2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR 2001)*, volume 1, 511–518.

Wang, Y.; Tan, T.; and Jain, A. K. 2003. Combining face and iris biometrics for identity verification. In *4th International Conference on Audio- and Video-Based Biometric Person Authentication (AVBPA'03)*, 805–813. Berlin, Heidelberg: Springer-Verlag.

Zhang, Z.; Wang, R.; Pan, K.; Li, S.; and Zhang, P. 2007. Fusion of near infrared face and iris biometrics. In Lee, S.-W., and Li, S., eds., *Advances in Biometrics*, volume 4642 of *Lecture Notes in Computer Science*. Springer Berlin / Heidelberg. 172–180.