

Overview of the Multiple Biometrics Grand Challenge

P. Jonathon Phillips¹, Patrick J. Flynn², J. Ross Beveridge³, W. Todd Scruggs⁴, Alice J. O’Toole⁵, David Bolme³, Kevin W. Bowyer², Bruce A. Draper³, Geof H. Givens⁶, Yui Man Lui³, Hassan Sahibzada¹, Joseph A. Scallan III⁷, and Samuel Weimer⁵

¹ National Institute of Standards and Technology, Gaithersburg, MD 20899, USA

² Computer Science & Engineering Department, U. of Notre Dame, Notre Dame, IN 46556, USA

³ Department of Computer Science, Colorado State U., Fort Collins, CO 80523, USA

⁴ SAIC, Arlington, VA 22203, USA

⁵ School of Behavioral and Brain Sciences, The U. of Texas at Dallas, Richardson, TX 75083, USA

⁶ Department of Statistics, Colorado State U., Fort Collins, CO 80523, USA

⁷ Schafer Corp., Arlington, VA 22203, USA

Abstract. The goal of the Multiple Biometrics Grand Challenge (MBGC) is to improve the performance of face and iris recognition technology from biometric samples acquired under unconstrained conditions. The MBGC is organized into three challenge problems. Each challenge problem relaxes the acquisition constraints in different directions. In the Portal Challenge Problem, the goal is to recognize people from near-infrared (NIR) and high definition (HD) video as they walk through a portal. Iris recognition can be performed from the NIR video and face recognition from the HD video. The availability of NIR and HD modalities allows for the development of fusion algorithms. The Still Face Challenge Problem has two primary goals. The first is to improve recognition performance from frontal and off angle still face images taken under uncontrolled indoor and outdoor lighting. The second is to improve recognition performance on still frontal face images that have been resized and compressed, as is required for electronic passports. In the Video Challenge Problem, the goal is to recognize people from video in unconstrained environments. The video is unconstrained in pose, illumination, and camera angle. All three challenge problems include a large data set, experiment descriptions, ground truth, and scoring code.

1 Introduction

Since 1993, progress in automatic face recognition has been assisted by a series of “challenge problems” and “independent evaluations” organized by the U.S. Government [1] [2] [3] [4] [5] [6]. The challenge problems were designed to advance the evolving state-of-the-art in the field of automatic face recognition.

The independent evaluations were designed to measure progress made in automatic face recognition. A common thread connecting these challenge problems and evaluations was benchmarking performance of frontal face recognition from still images taken under controlled illumination. Additionally, more difficult aspects of face recognition have been incorporated into most of these efforts. A fundamental result of the Face Recognition Vendor Test (FRVT) 2006 was its demonstration of significant progress, relative to earlier evaluations, on frontal face recognition from still images taken under controlled illumination [6].

Challenge problems and third party evaluation in iris recognition began in 2004 [5] [6] [7] [8] [9]. To better understand the state-of-the-art in iris recognition, Newton and Phillips [10] performed a meta-analysis on the ICE 2006, the ITIRT, and the Iris 06 evaluations. The finding of the meta-analysis was that despite the differences in the testing protocols, sensors, image quality, subject variability and failures to enroll and acquire, the performance results from all three evaluations were comparable⁸.

The Multiple Biometrics Grand Challenge (MBGC), currently underway, is aimed at examining a set of biometrics in a single evaluation program. An important part of this challenge is to relax the controlled imaging constraints used in previous tests. Again, the emphasis is to pose challenge problems that motivate improved recognition from biometric samples acquired under unconstrained conditions. In this paper, we provide an overview of the MBGC. The MBGC is organized into three challenge problems: the Portal Challenge Problem, the Still Face Challenge Problem, and the Video Challenge Problem. We begin with a brief statement of the goals of each of the three challenge problems. Next, we present a description of the data used in the challenge problems. Finally, we describe the protocol and experiments being applied in each challenge problem and sketch out the projected schedule for the MBGC.

2 Challenge Problems

In the Portal Challenge Problem, the goal is to recognize people from near-infrared (NIR) and high definition (HD) video as they walk through a portal. Iris recognition can be performed from the NIR video and face recognition from the HD video. Person recognition can be augmented beyond face or iris recognition alone by fusing the NIR and HD modalities. The Portal Challenge Problem is designed to examine identification from biometric data captured passively as a person passes through a portal. The Portal Challenge Problem for the first time makes NIR video clips available to researchers to support iris recognition research. Thus, solving the Portal Challenge Problem will require the development of a number of new technologies. First is the development of algorithms that can process the NIR video clips to locate ocular and iris regions and perform recognition from these regions. The second is the development of fusion algorithms to perform recognition from both the NIR and HD video clips. Because the biometric samples in the NIR and HD video sequences are acquired in

⁸ A survey of iris recognition can be found in Bowyer et al [11].

a portal environment, it is not guaranteed that high quality iris or face samples will be captured. The fusion algorithms need to be able to handle cases where an iris is not present or one of the samples is poorly imaged. The requirement for recognition from NIR video clips has the potential to extend recognition from the iris to the entire ocular region of the face.

The Still Face Challenge Problem builds on the Face Recognition Grand Challenge (FRGC) adding non-frontal face images to the task and increasing the size and scope of experiments with images acquired under uncontrolled illumination. A survey of work addressing face recognition from unconstrained environments can be found elsewhere [12] and provides some baseline estimates of the state-of-the-art for this problem. The increase in the scope of the Still Face Challenge Problem includes recognition from highly compressed images.

The Video Challenge Problem breaks new ground by being the first challenge problem to address face recognition from unconstrained video. It makes available information that is commonly included in surveillance video and so provides a look at the level of recognition that can be achieved in security applications where video is available.

The MBGC is being conducted in two parts with two versions of the challenge problems. One of the goals of the version 1 is to introduce participants to the challenge problems and the MBGC protocol. The goal of version 2 is to encourage the development of algorithms that can handle large datasets. The version 1 challenge problems were released starting in May 2008 and the version 2 challenge problems will be released starting in February 2009. This paper describes the MBGC as of January 2009 and primarily focuses on the version 1 challenge problems. The version 2 challenge problems builds on version 1. Version 2 will consist of the same three challenge problems with the same types of biometric samples. The primary difference between the two versions is that larger datasets will be included in version 2.

3 Data

The MBGC challenge problems consist of data collected at the U. of Notre Dame specifically to support the MBGC and a previously collected dataset at the U. of Texas at Dallas [13]. The data collected at the U. of Notre Dame are included in all three challenge problems. The data collected at the U. of Texas at Dallas are included in only the Video Challenge Problem.

The data in the Portal Challenge Problem consists of five recording types. Two of the recording types were collected from an LG 2200 sensor: still iris images and video iris sequences. The still iris images intentionally represent a broader range of quality than the LG 2200 sensor would normally acquire, see Phillips et al [5] [6] [9] for details. The LG 2200 iris video sequences were digitized from the National Television System Committee (NTSC) video signal from the sensor. The NTSC video signal produced by the LG 2200 was digitized by a DayStar XLR8 USB video digitizer attached to a Macintosh host system, and stored in uncompressed format. Initial processing edited each subject's original

video, which contained both eyes, into two clips, one containing the left eye only and one containing the right eye only. These clips were transcoded to MPEG-4 format with a high bit rate allowance, thus yielding near-lossless encoding.

The third recording is very high resolution NIR video acquired from a Sarnoff Iris on the Move (IoM) system [14], see Figure 1. The IoM system was designed to capture iris imagery as a person walks through a portal.

The fourth recording type was acquired with a High Definition (HD) video camera, The HD video camera was bore-sighted beside the NIR video cameras in the IoM chassis. This camera captured HD video as a subject walked through the portal, see Figure 1. The HD video sequences are distributed in portrait mode, and the dimension of the video are 1440 pixels high by 1080 pixels wide. Because the camera was in portrait mode, the video was compressed vertically and the interlace artifacts show up in the vertical direction. The fifth recording type was a medium resolution still digital frontal image.

The data for the Still Face Challenge Problem were collected with high resolution digital cameras, 4 and 6 mega-pixels. Images were collected with both controlled and uncontrolled illumination. The images with controlled illumination were collected in a studio environment with controlled lighting. The images with uncontrolled illumination were collected in hallways and outdoors. Both frontal and non-frontal face images were collected.

The data for the Video Challenge Problem were collected with HD and Standard Definition (SD) video cameras. The data were collected in hallways, atria, and outdoors with unconstrained pose and illumination, see Figure 2. In the Video Challenge Problem version 1, the data from the U. of Notre Dame consisted of two types of video sequences: walking and activity. In the walking sequences, the subjects walked towards the camera. In the activity sequences, the subjects performed an action that required subjects to look away from the camera. An example action is picking up or looking at a book. The data from U. of Texas at Dallas consist of two types of video sequences: walking and conversation. In the walking sequences the subjects walked towards the camera. In the conversation sequences, a camera is looking down on a conversation between two subjects. One subject’s back is to the camera and this subject is to be ignored in the experiments. The other subject’s face is facing the camera at an off angle. This is the subject to be recognized.

4 Portal Challenge Problem

The Portal Challenge Problem version 1 consists of five experiments. Each experiment was designed to address different aspects of the Portal Challenge Problem data. These aspects address single mode face and iris recognition, and multiple biometric recognition. The five experiments in the Portal Challenge Problem are briefly described below and summarized in Table 1.

Still Iris versus NIR: The target image is an iris still from an LG 2200 camera. The query image is a NIR video sequence. This experiment tests iris recognition.




Fig. 1. Example of imagery collected as a subject walks through the portal. The image in (a) is a full 2,000 by 2,000 pixel NIR frame acquired by the IoM and (b) is the left ocular region from the frame in (a). There are approximately 120 pixels across the iris in (a) and (b). The images in (c) are three frames from an HD video sequence of a subject walking through the portal.

Video Iris versus NIR: In this experiment both the target and query are video sequences. The target is an LG 2200 iris video and the query is a NIR video sequence. This experiment tests iris matching using video in both the target and query sets.

Still Face versus HD Video: The target is a still face images and the query images is an HD video sequence. This experiment tests facial recognition.

Still Iris and Still Face versus NIR and HD Video: This experiment combines the elements from the pervious two experiments (Video Iris vs. NIR, and Still Face vs. HD Video Face). The target images consist of a face still and an LG 2200 still iris. The query images consists of two video sequences. The first is an HD video sequences, and the second is a NIR video sequence. This experiment tests recognition from face and iris imagery.




Fig. 2. Example frames from video sequences in the Video Challenge Problem. The images in (a) and (b) are from the U. of Notre Dame collection. The image in (a) is from a video sequence of a subject walking towards the cameras in an atrium and (b) is a subject performing a non-frontal activity outdoors. The images in (c) and (d) are from the U. of Texas at Dallas collection. The image in (c) is from a video sequence of a subject walking towards the cameras in an atrium. The image in (d) is acquired from a video camera looking down on a conversation. In (d), the subject whose face can be seen is the subject to be recognized.

Video Iris and Still Face versus NIR and HD Video: This experiment is the same as previous experiment except that LG 2200 video iris replaces the LG 2200 still iris images.

5 Still Face Challenge Problem

The goal of the Still Face Challenge Problem is to improve face recognition performance on images taken under conditions that are representative of a broad range of real-world scenarios. One set of experiments uses high resolution still face images taken in both controlled and uncontrolled illumination environments. This set of experiments challenges participants to match ‘controlled to uncontrolled’ face images and to match ‘uncontrolled to uncontrolled’ face images. The target sets consist of frontal face images and the query sets consist of both frontal and non-frontal face images.

Table 1. A summary of the experiments in the Portal Challenge Problem version 1. The first column is the name of the experiment. The second and third columns respectively list the type of imagery in the target and query sets. The last column lists the primary biometric modality in the experiment.

Experiment	Target image	Query image	Problem
still iris vs NIR	still iris	NIR video	iris
video iris vs NIR	video Iris	NIR video	iris
still face vs HD	still face	HD video	face
still iris & still face vs NIR & HD	still iris & still face	NIR & HD video	multiple biometrics
video iris & still face vs NIR & HD	video iris & video face	NIR & HD video	multiple biometrics

A second set of experiments in the Still Face Challenge Problem is designed to model the size of face images to meet the International Civil Aviation Organization’s (ICAO) standard for electronic passports [15]. Two of the primary requirements of this ICAO standard are the size of the face and compression of the face image. The face is required to be frontal with a size range from 90 to 120 pixels between the centers of the eyes. The face image is required to be compressed to between 8 KBytes and 20 KBytes. The images can be compressed by either JPEG or JPEG2000. To measure the range of possible performance, two sets of images were generated from the frontal images in the Still Face Challenge Problem, see Figure 3. In the first set, the images were rescaled to 90 pixels between the eyes and compressed to 8 KBytes. In the second set, the images were rescaled to 120 pixels between the eyes and compressed to 20 KBytes. Both sets of images were compressed using JPEG2000.

6 Video Challenge Problem

The Video Challenge Problem is designed to encourage the development of face recognition from video taken in hallways, atria, and outdoors under unconstrained illumination, pose, and camera angle. Version 1 of the Video Challenge Problem examined the effects of HD and SD video on performance. Version 1 also examined the effects of matching walking to walking sequences, activity to activity sequences, and walking to activity sequences, see Figure 1. All comparisons in the Video Challenge Problem are between video sequences taken on different days. This is to avoid matching based on clothes. Human performance on recognizing people in video sequences will be measured from data in the Video Challenge Problem. Human performance experiments will be designed to allow for direct comparison of human and machine recognition performance on the same pairs of video sequences.



(a)



(b)

(c)

Fig. 3. Example of face images in the Still Face Challenge Problem. The image in (a) is an original still face image taken outdoors. The size of the image is 3008 pixels wide by 2000 pixels high, and the size of the face is 232 pixels between the centers of the eyes. The image in (b) is the original that has been rescaled to 120 pixels between the centers of the eyes, cropped, and compressed to 20Kbytes. The image in (c) is the original that has been rescaled to 90 pixels between the centers of the eyes, cropped, and compressed to 8Kbytes. The dimensions of the image in (b) are 408 pixels wide by 528 pixels high, and the image in (c) is 306 pixels wide by 396 pixels high.

7 Protocol

The MBGC version 1 data and challenge problems were made available to participants starting in May 2008⁹. Each challenge problem consisted of a data set, description of the experiments, ground truth, and code to score the experiments. Results from the MBGC version 1 were presented at the Second Multiple Biometrics Grand Challenge Organizational Workshop held on 5 December 2008. The MBGC version 2 data and challenge problems will be made available to participants starting in February 2009.

Participants could submit results for any experiment in one or more of the three challenge problems. Participants could submit results for multiple algorithms for an experiment. All groups that submitted results agreed in advance that their performance results could be publicly attributed to them.

The MBGC is open to academic institutions, research laboratories and to companies worldwide. Groups granted access to the MBGC data are not required to submit results for analysis. Participation in MBGC version 1 is not required for participation in MBGC version 2.

8 Conclusion

The MBGC challenge problems are designed to advance face and iris recognition performance in unconstrained scenarios of clear practical importance. For the Portal Challenge Problem, the goal is to recognize people as they walk through a portal. The challenge is to develop algorithms that can process both the iris and the ocular regions of the face, and to fuse both the NIR and HD video sequences. The goal of the Still Face Challenge Problem is to recognize people in uncontrolled still images, and to test the effects of compression on face recognition. The Video Challenge Problem is the first challenge problem to address recognition of people in unconstrained video. The MBGC and its participants are developing technology to support the next generation of biometric applications.

9 Acknowledgements

We acknowledge the support of the the Biometric Task Force, the Department of Homeland Security's Directorate for Science and Technology, the Intelligence Advanced Research Projects Activity (IARPA), the Federal Bureau of Investigation (FBI), and the Technical Support Working Group (TSWG). The identification of any commercial product or trade name does not imply endorsement or recommendation by the Colorado State U., the National Institute of Standards and Technology, SAIC, Schafer Corp., the U. of Notre Dame or the U. of Texas at Dallas.

⁹ Instructions for getting the MBGC challenge problems and participating in the MBGC can be found at <http://face.nist.gov/mbgc/>.

References

1. P. J. Phillips, H. Wechsler, J. Huang, and P. Rauss, "The FERET database and evaluation procedure for face-recognition algorithms," *Image and Vision Computing Journal*, vol. 16, no. 5, pp. 295–306, 1998.
2. P. J. Phillips, H. Moon, S. Rizvi, and P. Rauss, "The FERET evaluation methodology for face-recognition algorithms," *IEEE Trans. PAMI*, vol. 22, pp. 1090–1104, October 2000.
3. P. J. Phillips, P. J. Grother, R. J. Micheals, D. M. Blackburn, E. Tabassi, and J. M. Bone, "Face recognition vendor test 2002: Evaluation report," National Institute of Standards and Technology, Tech. Rep. NISTIR 6965, 2003, <http://www.frvt.org>.
4. P. J. Phillips, P. J. Flynn, T. Scruggs, K. W. Bowyer, J. Chang, K. Hoffman, J. Marques, J. Min, and W. Worek, "Overview of the face recognition grand challenge," in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2005, pp. 947–954.
5. P. J. Phillips, W. T. Scruggs, A. J. O’Toole, P. J. Flynn, K. W. Bowyer, C. L. Schott, and M. Sharpe, "FRVT 2006 and ICE 2006 large-scale results," National Institute of Standards and Technology, Tech. Rep. NISTIR 7408, 2007. [Online]. Available: <http://iris.nist.gov/>
6. —, "FRVT 2006 and ICE 2006 large-scale results," *IEEE Trans. PAMI*, (in Press).
7. Authenti-Corp, "Iris recognition study 2006 (IRIS06)," Authenti-Corp, Tech. Rep. version 0.40, March 2007. [Online]. Available: <http://www.authenti-corp.com/iris06/report/>
8. International Biometric Group, "Independent testing of iris recognition technology," International Biometric Group, Tech. Rep., May 2005. [Online]. Available: <http://www.ibgweb.com/reports/public/ITIRT.html>
9. P. Phillips, K. W. Bowyer, P. J. Flynn, X. Liu, and W. T. Scruggs, "The Iris Challenge Evaluation 2005," in *Second IEEE International Conference on Biometrics: Theory, Applications, and Systems*, 2008.
10. E. M. Newton and P. J. Phillips, "Meta-analysis of third-party evaluations of iris recognition," *IEEE Trans. on SMC-A*, vol. 39, no. 1, pp. 4 – 11, 2009.
11. K. W. Bowyer, K. Hollingsworth, and P. J. Flynn, "Image understanding for iris biometrics: a survey," *Computer Vision and Image Understanding*, vol. 110, no. 2, pp. 281–307, 2008.
12. W. Zhao, R. Chellappa, P. J. Phillips, and A. Rosenfeld, "Face recognition: A literature survey," *ACM Computing Surveys*, vol. 35, pp. 399–458, 2003.
13. A. J. O’Toole, J. Harms, S. L. Snow, D. R. Hurst, M. R. Pappas, J. H. Ayyad, and H. Abdi, "A video database of moving faces and people," *IEEE Trans. PAMI*, vol. 27, pp. 812–816, 2005.
14. J. R. Matey, O. Naroditsky, K. Hanna, R. Kolcynski, D. J. Lofancono, S. Mangru, M. Tinker, T. M. Zappa, and W. Y. Zhao, "Iris on the move: Acquisition of images for iris recognition in less constrained environments," *Proceedings of the IEEE*, vol. 94, pp. 1936–1947, 2006.
15. ICAO New Technologies Working Group, "ICAO Document 9303," Tech. Rep. [Online]. Available: <http://www2.icao.int/en/MRTD/Pages/Doc9393.aspx>