Hyperspectral Face Database

Louis J. Denes, Peter Metes, and Yanxi Liu

CMU-RI-TR-02-25

October 2002

Robotics Institute
Carnegie Mellon University
Pittsburgh, Pennsylvania 15213

Abstract

In October 2001 we started to collect hyperspectral face images covering the spectral range from 450nm to 1100 nm. To date the database contains 54 diverse faces at multiple sessions over a period of about two months. The data was obtained using a prototype (limited performance) CMU-developed spectropolarimetric camera. We present some experimental results on the spectral face data. The unprocessed HID data is available upon permission. Recently we have developed a more sensitive next generation camera (Gen 2) that expands our data collection capabilities.

Contents

Hyperspectral Face Database

1 Introduction

Human identification from facial features has been studied primarily using imagery from visible video cameras. Several examples of such imagery are available in recent CMU reports [1,2]. Reference [3] is a recent overview of issues involving face recognition. Spectropolarimetric imaging is an innovative emerging technology in the market. The AOTF (Acousto-Optic-Tunable-Filter) is a birefringent crystal capable of rapid and precise wavelength selection. It may be used as the heart of a spectrometer or monochromator for applications in the UV through mid-IR. Advantages over a filter wheel or grating monochromator include high resolution, high speed, random or sequential wavelengths access, no moving parts, compact size, and imaging capabilities. When combined with the advanced electronics, software, and optics the AOTF becomes a unique scientific instrument for many applications. The CMU AOTF System operates in visible to NIR from 450nm to 1100nm. Other research groups led by Glenn Healey (www.cvl.uci.edu) at UC at Irvine and Larry Wolff and colleagues at Equinox (www.equinoxsensors.com/news.html) have also addressed various aspects of spectral imaging directed at facial recognition [4-7]. Moritz Stoerring et al. investigated the detailed spectral characteristics of the skin in the visible-near IR that also includes racial color aspects. Their findings are available at the web site (www.vision.auc.dk/~mst/Publications/sirs99html/paper final.html).

This paper describes the capture setup and illumination geometry, presents the calibration methodology for the spectropolarimetric camera, present examples of the face collection database, featuring the database used to test spectral face asymmetry for HID. The report concludes with presentation of the performance improvement of a newly built Gen 2 camera that overcomes deficiencies observed using the Gen 1 camera.

2 Capture Setup

Figure 1 pictures the studio setup used in the hyperspectral face data collection. The view shows the relative positions of the camera, the lights and the human subject. For illumination we used three identical lamps with 600W halogen bulbs. We need this level of power because of relative lack of sensitivity of the system (only about 5-10% of light is useful).

We use a head rest stop so that all subject's heads are firmly positioned and located at nearly the same height. We instructed each subject to maintain a still pose during the data collection session. However, because of the brightness of the studio lamps some eye blink was inevitable.

Figure 2 shows in planar view the overall studio layout for the hyperspectral face data collection. The main components include three illumination sources, the hyperspectral camera, control and data collection computer as depicted in planar view. The lamps are at angles of –45, 0, and +45 degrees with respect to the face subject. All operations are under computer control. The three lamps can be turned on individually or together.

The spectral face identification database is obtained under defined illumination. Figure 2 shows the light position and sequence for each set of data.

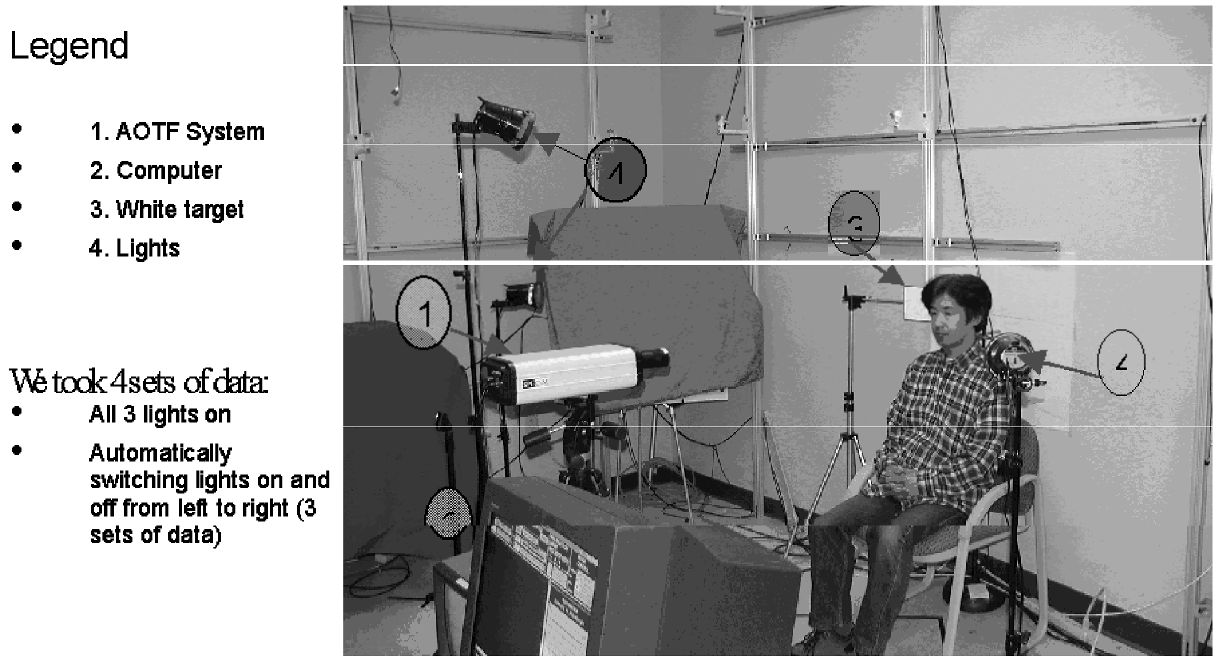

Figure 1. Setup of the studio for data collection

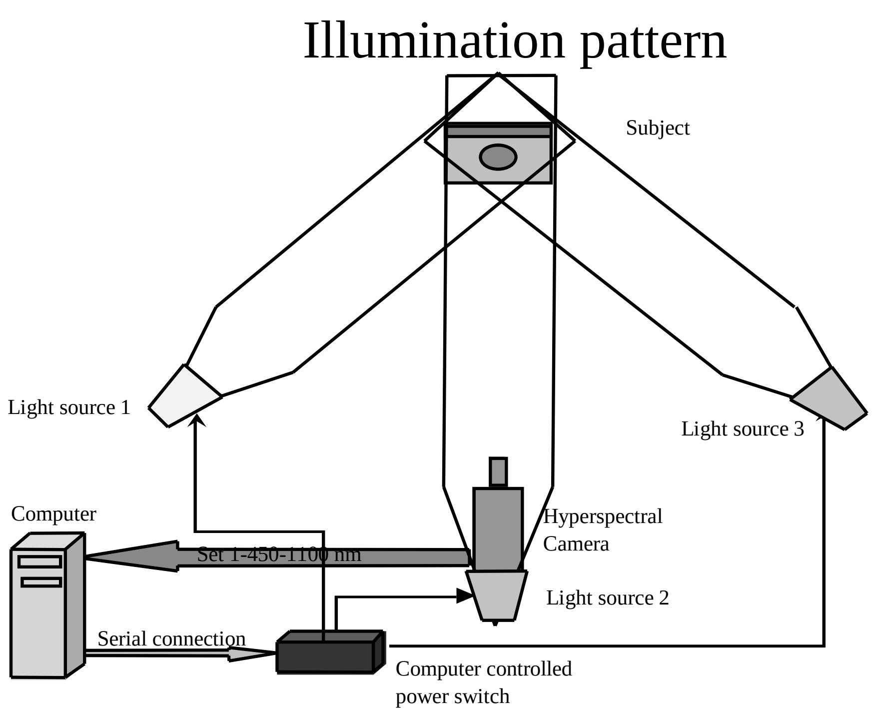

Figure 2. Studio layout used in the hyperspectral face data collection.

We took face data on 54 subjects between 10/18/01 and 12/4/01. Multiple sessions (up to five) were taken on most of the subjects over a period of several weeks. Each session contains 20 Mbytes of data consisting of four multispectral sequences – illumination from 45 right, center, 45 left and total. The hyperspectral range is 450 to 1100 nm in steps of 10 nm (65 spectral bands). Without specific approval, this database is restricted to personnel involved with CMU's DARPA HID program. **The hyperspectral database is located at: /hid/data/hyperspectral.**

Hyperspectral camera

Figure 3 pictures the hyperspectral camera (i.e., spectro-polarimetric imaging camera) outside its enclosure along with its nominal specifications. Spectrally the camera covers 450 to 1100 nm with a spectral band pass of 10 nm at 600 nm. We did not use the polarimetric capabilities of this camera in this study. The camera control software is written in Visual Basic, a language that suffices for the medium resolution (640 x 480) analog camera using an external frame grabber. In this situation, the desktop computer controls the spectral filtering hardware while the frame grabber firmware/software handles the frame acquisition and presentation functions to the desktop computer. The display interface of the Visual Basic code is shown below (Figure 4).

Spectro-Polarimetric Imaging Camera

Specifications

AO material	TeO_2
Spectral range	450-1100 nm
Resolution	10 nm @600nm
AO efficiency	>80%
RF range	25-70 MHz
Retarder range	400-1800 nm
IFOV	~7deg
(Optics Adjustable)	
RF power	<1 W
AOTF aperture	15 x 15 mm
AO Interaction	15 mm
Crystal length	26.5 mm
Min. Illumination	CCD Camera dependent

Figure 3. Spectropolarimetric camera (uncovered) and performance features

Figure 4. Spectropolarimetric Desktop control

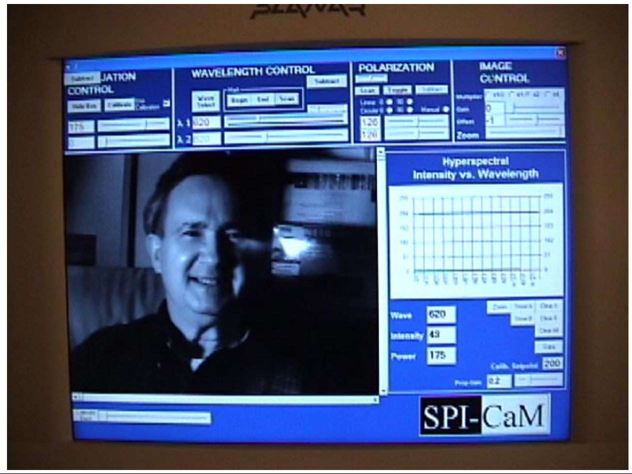

3 Hyperspectral system calibration

Two key properties in hyperspectral imaging are its spectral resolution and its reflectance response from a white reference. Ideally, we seek a spectrally flat response from a 'white' calibration target. The spectral resolution to a 'line' source (i.e., a HeNe laser beam) is presented in Figure 5a. The spread at FWHM is about 10 nm in accordance with the AOTF design for that wavelength. Shown in Figure 5b are three calibrated spectral responses to reflectance broadband incandescent lamps at angles of –45, 0 and +45 degrees to a spectralon target placed at the position of the subject's head. Note that with each lamp, the camera's response is quite flat over the 450 to 1100 nm spectral range covered.

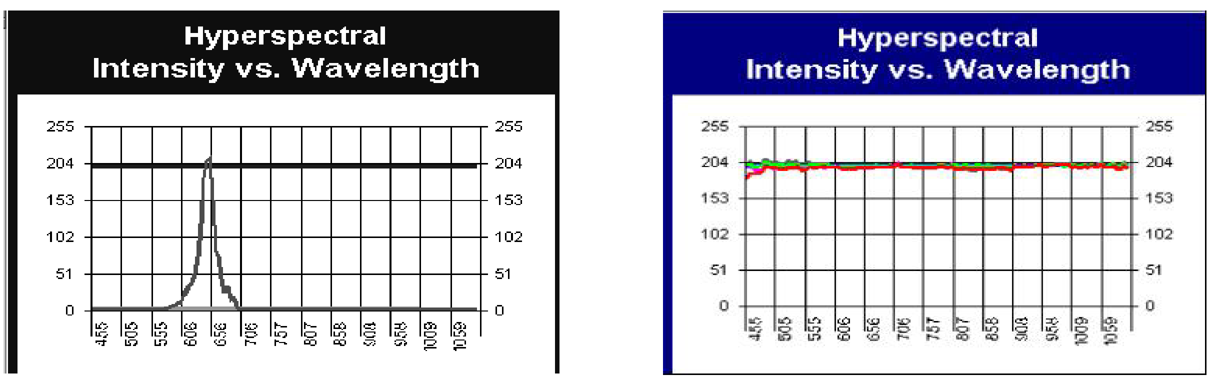

a) HeNe Laser beam spot

b) Studio lamps calibrated at 0, +45 and –45 degrees to subject

Figure 5. Spectral response to a line source and to broadband studio lamps.

CMU AOTF Spectropolarimetric Imaging System design employs a two level, black and white, calibration scheme. The black calibration level is performed only once to compensate for camera noise. The white calibration is performed as necessary on the fly. White calibration remained fairly constant during this test sequence since the illumination sources were reasonably stable. Calibrating the system with a white target will give a linear response through the analyzed spectrum for different illuminations as presented in Fig. 5b. As another check, we use a Gretag Macbeth ColorChecker to verify the linearity of spectral response of the spectropolarimetric camera and its response to ‘standard’ Macbeth colors. The figure in the upper right hand corner of Figure 8 compares the measured spectral response to the Macbeth gray level standard. The linearity is seen to be quite good only falling off at the ends where the detector response is weak. Figure 6 is a survey of the relative positions of the light sources and the hyperspectral camera with respect to the face subject. All components including the centroid of the subject’s head is at 51.5 in above the studio floor. The exception is the second light source, which is elevated to 72 inches.

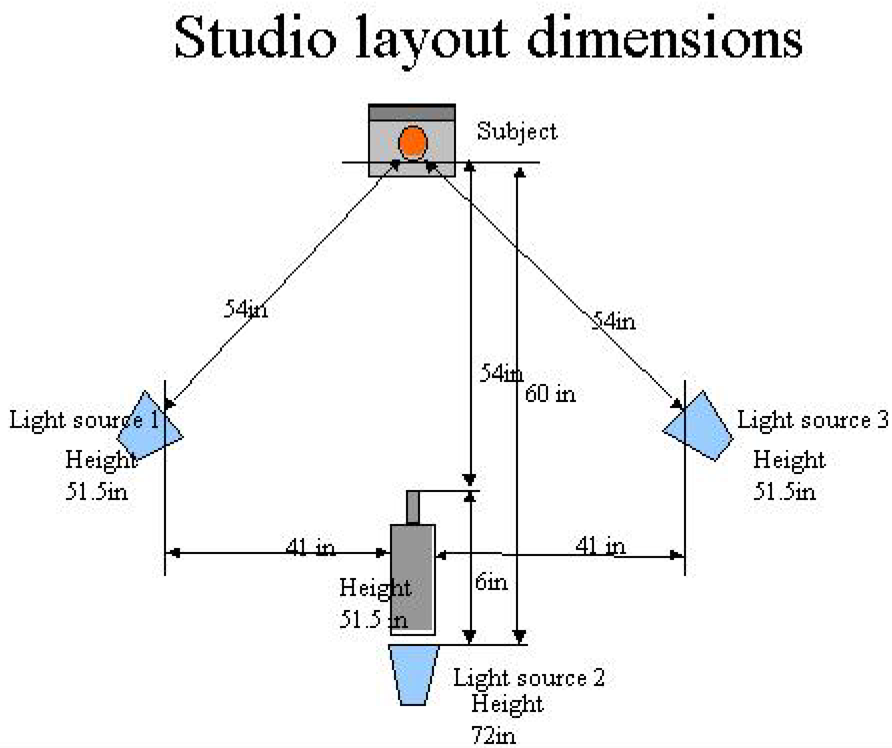

Figure 6. Studio layout dimensions.

We placed a checker board panel consisting of an array of 1 inch squares at the position of the subject’s head as a means for geometric calibration of the studio scene. The photos presented in Figure 7 are taken by positioning a digital camera at the respective locations of the lamps and hyperspectral camera as indicated.

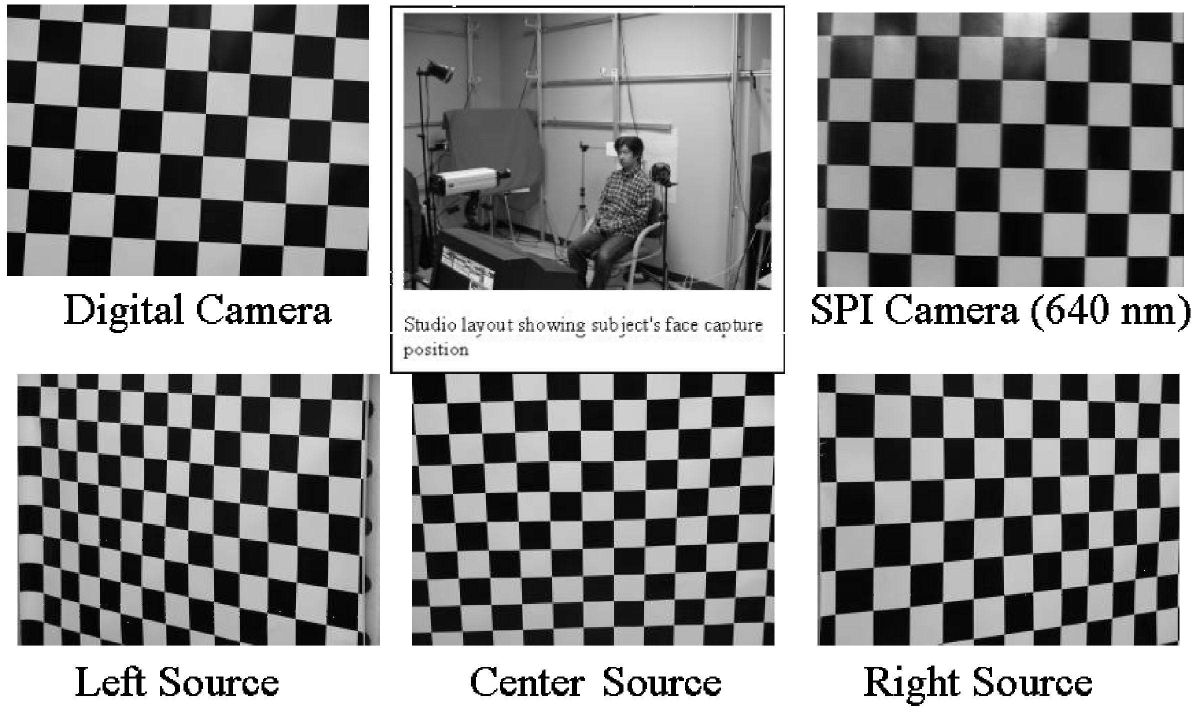

Figure 7. One inch squares 'checker board' used in the geometric calibration of the studio scene viewed from light sources and from the hyperspectral camera lens.

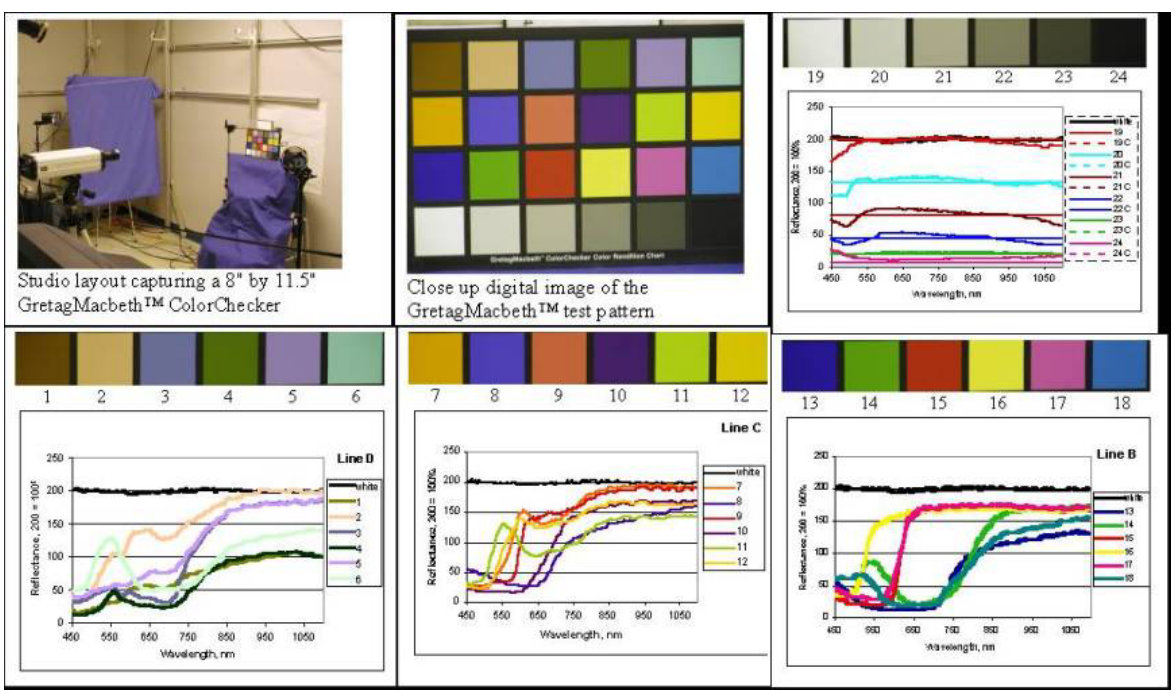

Figure 8. Gretag Macbeth ColorChecker Calibration.

4 Hyperspectral face identification database

Data files for all of the subjects and their sessions are located within the CMU/CSC system: **/hid/data/hyperspectral/**. Each folder has a number, assigned for each human subject. Depending on the number of sessions attended by subject there are subfolders for each meeting. Each session has four data (sub) folders. (Ex for subject 04521: Folder name: 04521, subfolders: sessionN, sessNCenter, sessNLeft and sessNRight where N is a number). Images in folder SessionN are with all lights on, in SessNCenter only center light was on, in SessNLeft only left light was on and in SessNRight only right light was on.

In this section we give illustrative examples of the hyperspectral face data that has been recorded. In all, we took data on 54 subjects between the dates of 10/18/01 and 12/4/01. At that point (12/4/01), we suspended data collection owing to feedback from test subjects that the photo lamps caused some residual eye irritation of a few of the subjects during and shortly after exposure. We sensed also that some accumulative increased eye sensitivity could occur due to the photo lamps intensity. In response, we set out to construct the Gen 2 spectropolarimetric imager that will operate effectively with lower lighting conditions with good signal to noise throughout the spectral range.

The session breakdown of the 54 subjects is the following:

28 subjects – 5 sessions
4 subjects – 4 sessions
3 subjects – 3 sessions
7 subjects – 2 sessions
12 subjects – 1 session

Each session (20 Mbytes) consisted of <u>four</u> multispectral sequences – illumination from: 1) 45 right, 2) center, 3) 45 left and 4) total. The range covered is 450 to 1100 nm in steps of 10 nm (65 spectral bands).

Figure 9 displays the effect on face pose using the four lighting scenarios covering five (550, 650, 750, 850 and 1000 nm) hyperspectral wavelengths. Notice that the ‘white’ calibration target at the left picture edge has about the same intensity while the facial pose is much darker and spatially noisy at the shorter (550 nm) and the longer (1000 nm) filtered exposures.

The next example (Figure 10) displays a facial pose every 50 nm from 450 nm to 1050 nm. Again we observe the darkened, spatial noisy exposures at the low and high end of the spectrally filtered views. The basic dilemma is that the exposure of our CCD camera is set at its fixed maximum value of 1/60 sec. Lighting and CCD has its maximum response in the around 650 nm. At shorter and longer wavelengths there are not enough photons to produce noise-free images. The lamp intensities (600 W) are already at near the upper end of commercial studio lighting. In the following section, where we have begun a face identification analysis, we find that darkened noisy images are not sufficient to provide adequate discrimination using current face asymmetry algorithms.

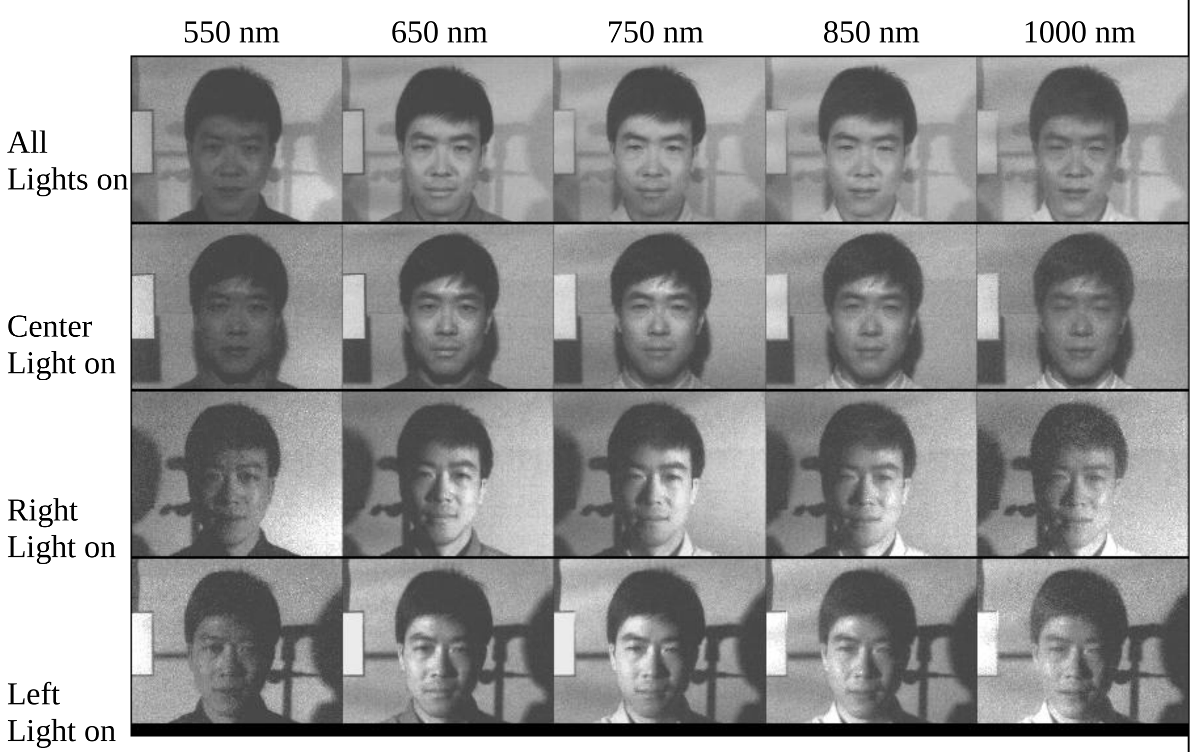

Figure 9. Sample Hyperspectral Face Data showing lighting effects.

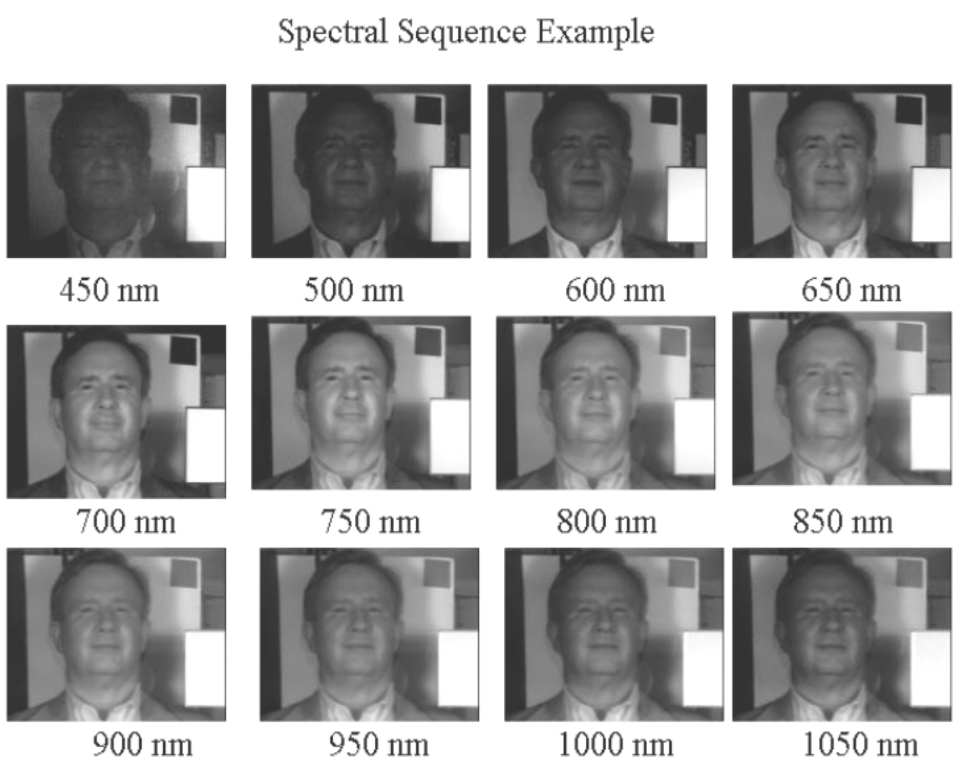

Figure 10. Hyperspectral face images covering the 450 to 1100 nm spectral range.

5 Database used to test spectral face asymmetry for HID

We constructed a database of 18 subjects to test spectral face asymmetry for HID. Figure 11 presents at the top the 18 training faces obtained at a filtered wavelength of 600 nm from the first session. The bottom set of frames presents the test face at a filtered wavelength of 700 nm taken during the second (days later) session. From these we generate the corresponding eigenfaces as pictured between the two test samples.

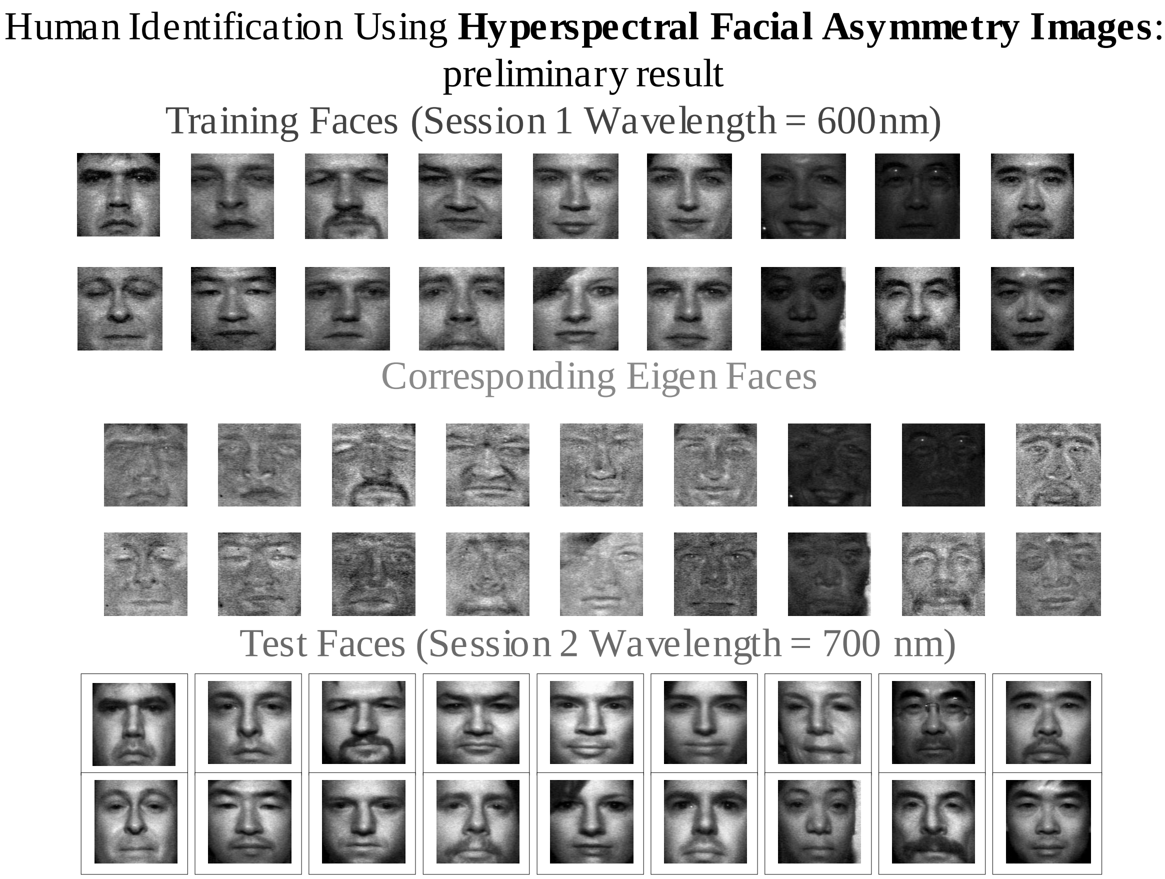

Figure 11. Database used for testing hyperspectral asymmetry faces.

Figures 12a and 12b present experimental results analyzing the hyperspectral face data to obtain face identification. Plotted is the number of unrecognized faces. Figure 12a has the same lighting conditions between sessions 1 and 2. Using 600 nm for the training set, recognition approaches 100 % from about 500 nm through 1100 nm for this small 18-subject database. (Picture quality is poor below 500 nm.) Recognition becomes increasingly worse using 700 nm and 800 nm for the training sets; still recognition is in the 80 – 90 % range. Figure 12b compares the faces of the two sessions under different lighting conditions. (Session 1 under centered illumination, Session 2 under left side illumination). We see that that the ability of facial recognition is reduced to less than 50% throughout the entire 450 to 1100 nm spectral range. Better face recognition clearly requires higher definition through a more sensitive, low noise camera or through higher levels of illumination. Information on the details of the algorithms used in this analysis is available fromY. Liu et al. references [8-12].

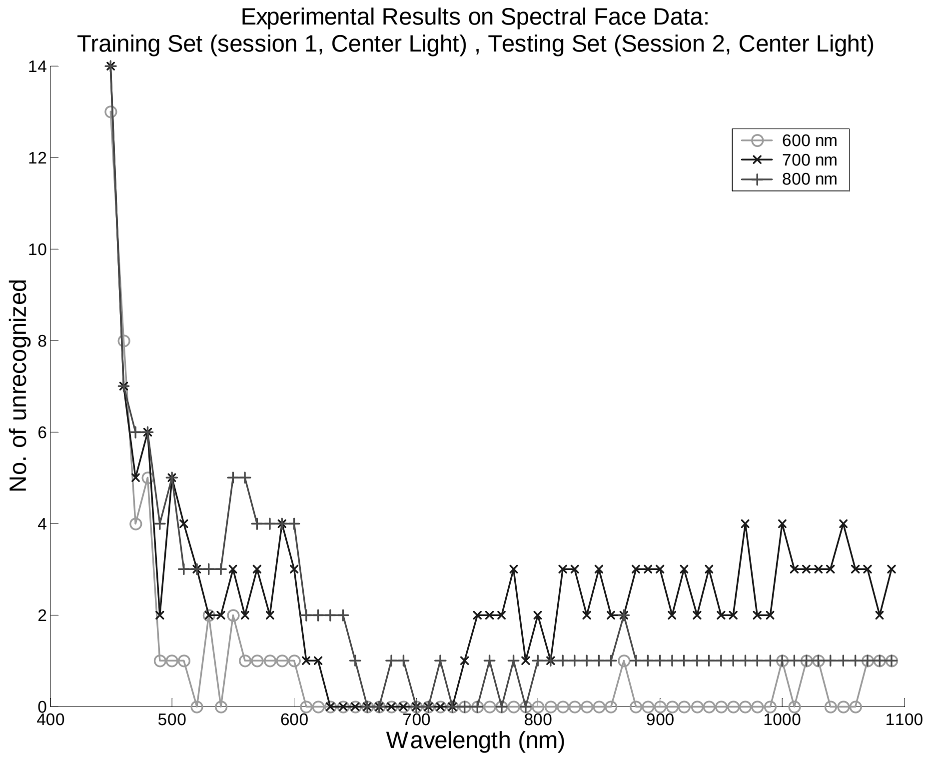

Figure 12a. Experimental Results on Spectral Face Data:
Training Set (session 1, Center Light). Testing Set (Session 2, Center Light)

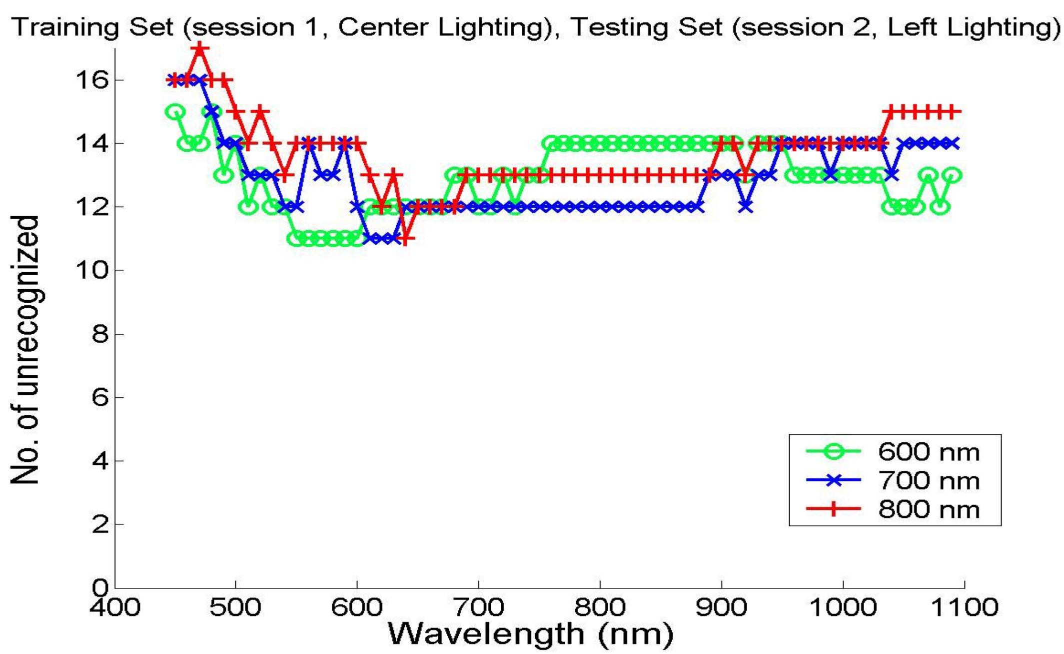

Figure 12b. Results on Spectral Face Data using different illumination:
Training Set (session 1, Center Light). Testing Set (Session 2, Center Light)

6 Future work

We plan to continue our data collection activities into the next period. A next generation (Gen 2) hyperspectral camera is under construction that will allow coverage of the 450 to 1100 nm spectral range with a better overall signal to noise and, more importantly, with lower external lighting requirements. This may come at the expense of a longer data collection sequence. We directed much of our efforts towards the development of the software for the next generation (Gen 2) hyperspectral imager. Of particular consequence for analog CCD cameras is a general absence of an electronic exposure control and its limited sensitivity especially below 600 nm. A second drawback is the need of a frame grabber with a requisite desktop size computer.

Recent high performance cameras utilize Sony's EXview HAD CCD™ sensor technology. These sensors have extended near infrared sensitivity in addition to greater overall sensitivity, broader dynamic range and better smear performance. Camera manufacturers have added features such as cooling, digital interfaces (e.g. IEEE 1394 'Firewire') and electronic shutter exposure that permit long exposures under computer control. We have selected a recently introduced camera, the QImaging Retiga 1350 EX.

This state-of-the-art digital camera combines the performance attributes of a cooled (0.1 e-/p/sec) CCD and an IEEE 1394 'Firewire' interface that provides high sensitivity, good dynamic range and reasonably high frame rate (low smear) display and data recording.

We incorporated a different calibration strategy. In the Gen 1 spectropolarimetric imager, we lacked the ability to electronically control the CCD camera exposure. In the previous calibration sequence, we first adjust the AOTF's RF drive power to best approximate a flat response by the imager to a calibrated 'white' target. The next sequence we adjusted the frame grabber's gain and offset to best achieve spectral flatness. The final step is to provide a mathematical correction that achieves the true spectral flatness. There are two unsatisfactory aspects to this calibration strategy. 1) The RF drive power is non-linear and can result in non-linear spectral response to lighting changes. 2) The most undesirable aspect is that although the flat white level is attained, it comes at the expense of a changing signal to noise (s/n) across the spectral imaging scan. This results in particularly grainy frames at values below 600 nm. Multiple exposures and averaging for spectral elements below 600 nm ought to work but at the cost of increased data collection time.

Use of a digital camera with a fast Firewire interface provides a solid calibration strategy. The RF drive power, camera gain and offset are predetermined and held constant. Exposure control is the only variable needed for flat spectral response to a calibrated 'white' target. Provided that the data is taken in a reasonably short time period, the frame s/n ratio will be constant resulting in the ability to analyze each spectral frame on equal terms. We have recently implemented these algorithms. Initial findings on the performance of this Gen 2 camera are quite encouraging. Direct comparison of the Gen 1 and Gen 2 camera images are displayed in Figure 10 a and 10 b. Gen 1 was calibrated using the original calibration routine; Gen 2 was calibrated using the exposure time procedure. The cameras were placed side by side and the spectral facial images were taken under the same illumination conditions at the same time. The images labeled (e.g., 500) correspond to Gen 1 while the images labeled e.g., 500Q correspond to Gen 2.

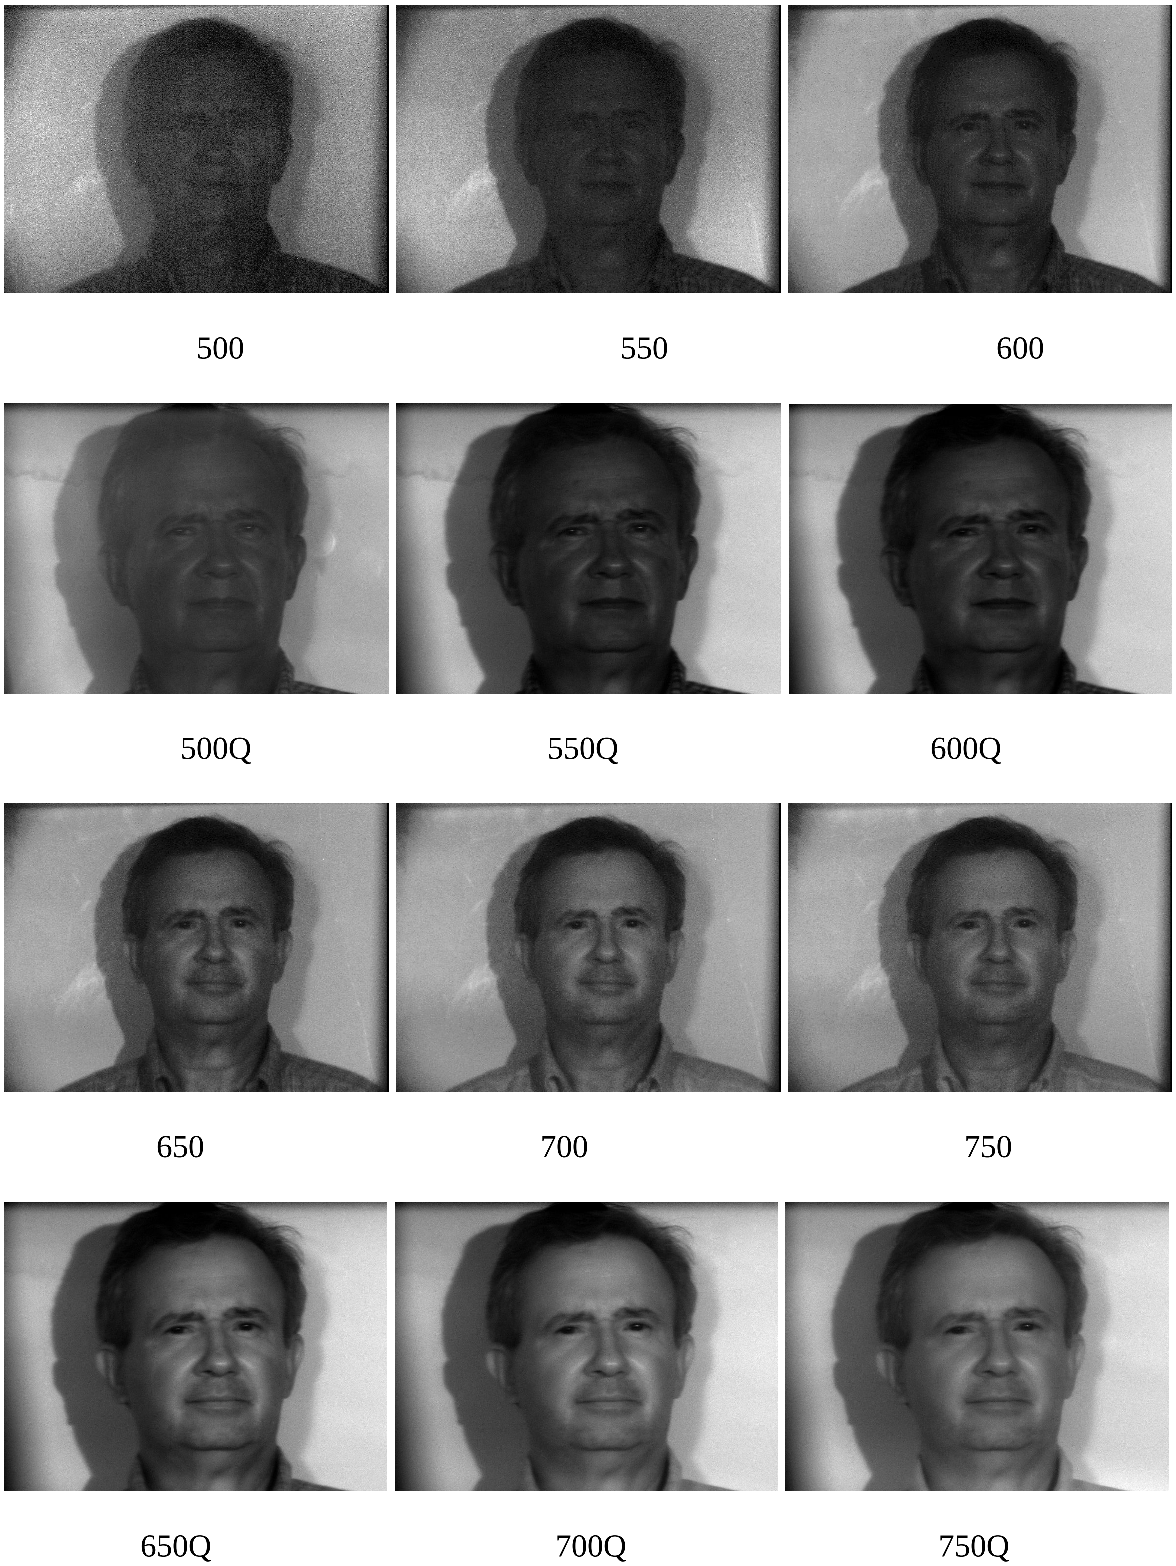

Figure 13a. Comparison of the spectral image quality of Gen 2 to Gen 1 cameras (500 to 750 nm).

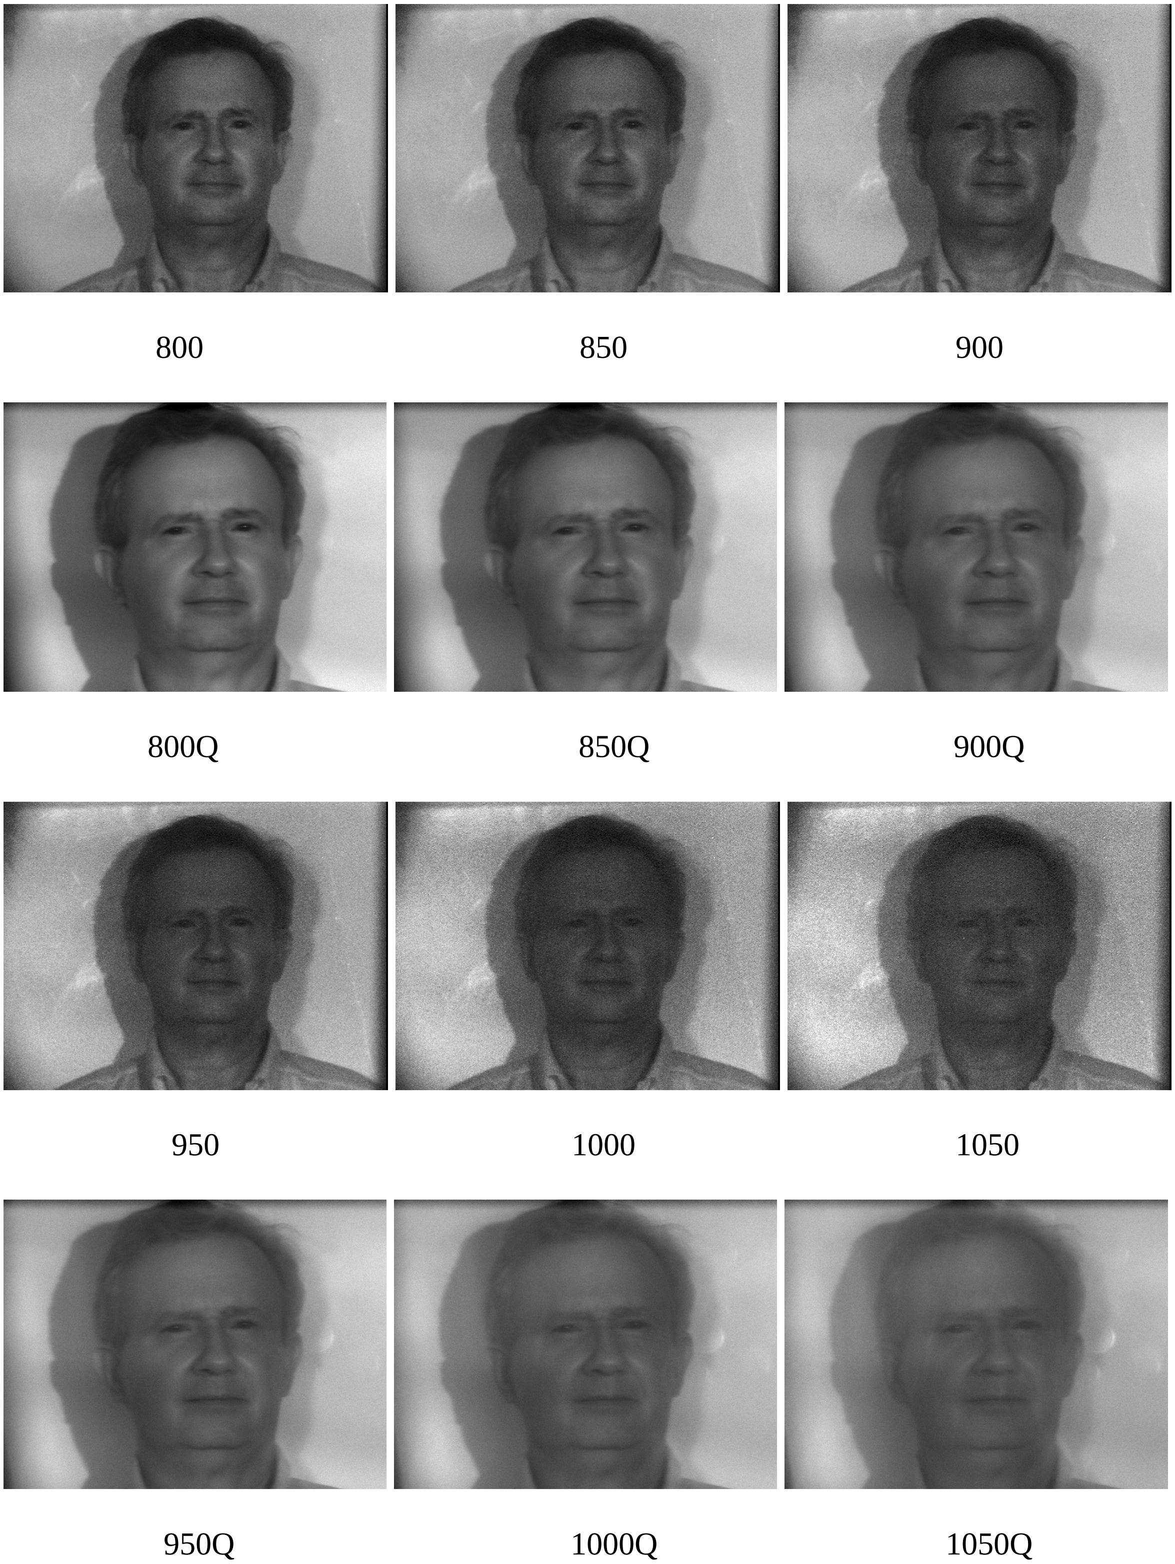

Figure 13b. Comparison of the spectral image quality of Gen 2 to Gen 1 cameras (800 to 1050 nm).

The picture quality of the Gen 2 camera, while improved using exposure balance, is still clearly inadequate at the short and long ends of the camera's spectral range. We find that the source of the problem is attributable to 'stray' light that leaks through the aperture of the AOTF and impinges on the CCD. This is broad band light which emperically we find to be closely proportional to the frame exposure. We have subsequently modified our calibration procedure to account of this exposure dependence contribution. Figure 14 presents an example using this exposure-modified calibration methodology. Note that the image quality is improved at both the short and long ends of the spectral range. There are opportunities for further improvement. The numerous spots seen in the images at longer exposures (500 nm, and 950 and beyond) are minute imperfections (i.e., dust or scratches) on the various optical surfaces within the optical train. The overriding limitation, however, lies in the high sensitivity region of spectral response of the current QImaging CCD camera that is not broad enough to cover the intended region of spectral interest. We expect to remedy this deficiency in the very near term with the next generation CCD cameras that are now being introduced.

In conclusion, we intend to extend our data collection efforts to involve:

1) More hyperspectral image collection of face subjects for identification evaluations,
2) Spectropolarimetric image collection for dehazing of atmospheric obscuration.

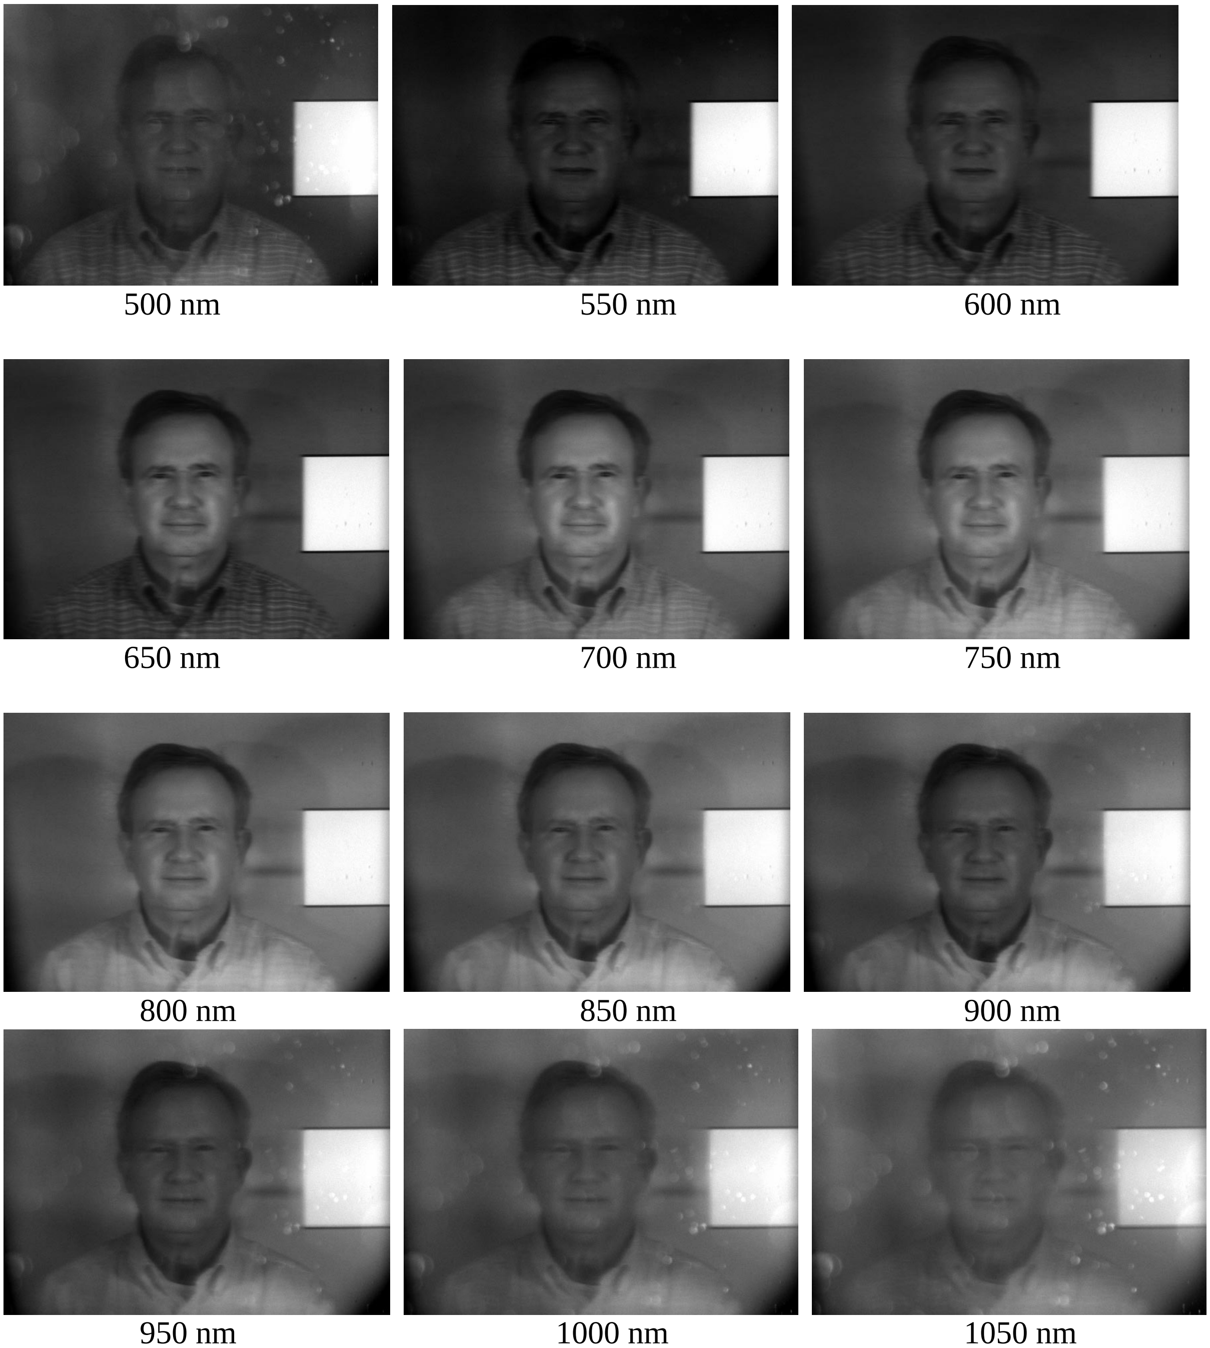

Figure 14. Hyperspectral facial images obtained using exposure-corrected calibration.

7 References

1. T. Sim, S. Baker, and M. Bsat " The CMU Pose, Illumination, and Expression (PIE) Database", Proceedings of the IEEE International Conference on Automatic Face and Gesture Recognition, May 2002. Also: Terence Sim, Simon Baker, and Maan Bsat " The CMU Pose, Illumination, and Expression (PIE) Database of Human Faces", Carnegie Mellon Report, CMU-RI-TR-01-02.
2. Ralph Gross and Jianbo Shi, "The CMU Motion of Body (MoBo) Database", Carnegie Mellon Report, CMU-RI-TR-01-18.
3. Warren Webb, "Friend or Foe - Face recognition goes to war", EDN (www.ednmag.com), pp. 32-35, 12/20/01.
4. Socolinsky, D., Wolff, L., Neuheisel, J., and Eveland, C. *Illumination Invariant Face Recognition Using Thermal Infrared Imagery*, Computer Vision and Pattern Recognition, Kauai, December 2001.
5. Eveland, C., Socolinsky, D., and Wolff, L. *Tracking Human Faces in Infrared Video*, CVPR Workshop on Computer Vision Beyond the Visible Spectrum, Kauai, December 2001.
6. Wolff, L., Socolinsky, D., and Eveland, C. *Quantitative Measurement of Illumination Invariance for Face Recognition Using Thermal Infrared Imagery*, CVPR Workshop on Computer Vision Beyond the Visible Spectrum, Kauai, December 2001.
7. Socolinsky, D. and Selinger A. *A Comparative Analysis of Face Recognition Performance with Visible and Thermal Infrared Imagery*. To appear ICPR '02, Quebec, August 2002.
8. Y. Liu, K. Schmidt, J. Cohn, and R.L. Weaver, "Facial Asymmetry Quantification for Expression Invariant Human Identification", Proceedings of the 5th International Conference on Automatic Face and Gesture Recognition (FG'02), Washington D.C., May 20-21, 2002.
9. Liu, Y., J.F. Cohn, K.L. Schmidt and Mitra, S. ``Facial Asymmetry Quantification for Expression Invariant Human Identification", to appear at Computer Vision and Image Understanding Journal, 2002.
10. Y. Liu, K. Schmidt, J. Cohn, and R.L. Weaver, "Human facial asymmetry for expression-invariant facial identification", Proceedings of the Fifth IEEE International Conference on Automatic Face and Gesture Recognition (FG'02), May, 2002.
11. Wachtman, G.S, Liu Y., Zhao T., Cohn J.F., Schmidt, K.L., Henkelmann,T.C, T.C., VanSwearingen J.M., Manders, E.K. "Measurement of Asymmetry in Persons with Facial Paralysis", Combined Annual Meeting of the Robert H. Ivy and Ohio Valley Societies of Plastic and Reconstructive Surgeons, June 8, 2002. (Best paper overall and first prize in Clinical Science category).
12. Y. Liu, R.L. Weaver, K. Schmidt, N. Serban, and J. Cohn, ``Facial Asymmetry: A New Biometric", CMU RI tech report CMU-RI-TR-01-23, 2001.

8 Acknowledgement

We thank Takeo Kanade and Jianbo Shi for their support and encouragement.
We also acknowledge the assistance of Rob Engel, John Ernsthausen and Daniel Ionescu in the development of the Gen 2 spectropolarimetric camera. Financial support was provided primarily by the U.S. Office of Naval Research (ONR) under contract N00014-00-1-0915.