Y/NSP-553

Laptop Computer-Based
Facial Recognition System Assessment

Ron Cain
Advanced Computing Technologies
BWXT Y-12, L.L.C.
Oak Ridge, Tennessee 37831
(865) 574-4042
cainra@y12.doe.gov

George B. Singleton
National Security Program Office
BWXT Y-12, L.L.C.
Oak Ridge, Tennessee 37831
(865) 241-3854
singletongb@y12.doe.gov

March 2001

prepared by
NATIONAL SECURITY PROGRAM OFFICE
Y-12 NATIONAL SECURITY COMPLEX
OAK RIDGE, TENNESSEE 37831-8260
managed by
BWXT Y-12, L.L.C.
for the
U.S. DEPARTMENT OF ENERGY
under contract DE-AC05-00OR22800

Y/NSP-553

Laptop Computer-Based Facial Recognition System Assessment

Ron Cain
Advanced Computing Technologies
BWXT Y-12, L.L.C.
Oak Ridge, Tennessee 37831

George B. Singleton
National Security Program Office
BWXT Y-12, L.L.C.
Oak Ridge, Tennessee 37831

March 2001

prepared by
NATIONAL SECURITY PROGRAM OFFICE
Y-12 NATIONAL SECURITY COMPLEX
OAK RIDGE, TENNESSEE 37831-8260
managed by
BWXT Y-12, L.L.C.
for the
U.S. DEPARTMENT OF ENERGY
under contract DE-AC05-00OR22800

CONTENTS

1. INTRODUCTION

The objective of this project was to assess the performance of the leading commercial-off-the-shelf (COTS) facial recognition software package when used as a laptop application. We performed the assessment to determine the system's usefulness for enrolling facial images in a database from remote locations and conducting real-time searches against a database of previously enrolled images. The assessment involved creating a database of 40 images and conducting 2 series of tests to determine the product's ability to recognize and match subject faces under varying conditions. This report describes the test results and includes a description of the factors affecting the results.

After an extensive market survey, we selected Visionics' FaceIt® software package for evaluation and a review of the Facial Recognition Vendor Test 2000 (FRVT 2000). This test was co-sponsored by the U.S. Department of Defense (DOD) Counterdrug Technology Development Program Office, the National Institute of Justice, and the Defense Advanced Research Projects Agency (DARPA). Administered in May–June 2000, the FRVT 2000 assessed the capabilities of facial recognition systems that were currently available for purchase on the U.S. market. Our selection of this Visionics product does not indicate that it is the "best" facial recognition software package for all uses. It was the most appropriate package based on the specific applications and requirements for this specific application.

In this assessment, the system configuration was evaluated for effectiveness in identifying individuals by searching for facial images captured from video displays against those stored in a facial image database. An additional criterion was that the system be capable of operating discretely. For this application, an operational facial recognition system would consist of one central computer hosting the master image database with multiple standalone systems configured with duplicates of the master operating in remote locations. Remote users could perform real-time searches where network connectivity is not available. As images are enrolled at the remote locations, periodic database synchronization is necessary (see Fig. 1).

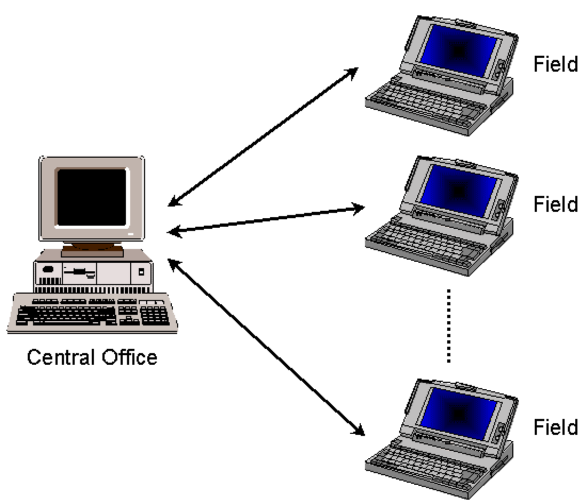

Fig. 1. Remote Operation Configuration

Because database synchronization is mainly a logistical consideration, our assessment focused primarily on system efficacy with respect to its ability to capture, store, manage, and respond to facial images. Therefore, the test setting did not include a central office component. The software was installed directly onto a laptop computer, and all capture and search processes were performed on the laptop system.

2. PRODUCT TESTED

Visionics software was selected for evaluation. Visionics Corporation (www.visionics.com), co-headquartered in New Jersey and Minnesota, develops and deploys facial recognition technology that allows computers to rapidly and accurately recognize faces. Visionics' software is promoted for authentication (banking, smart cards, access control, and border control); identification solutions (voter registration, passports, drivers licenses, and employee identification cards); criminal justice systems (mug-shot/booking systems and post-event analysis); and human identification at a distance (surveillance, human traffic control, and friend or foe).

Visionics software was chosen because of its placement in DOD Facial Recognition Technology (FERET) tests. The DOD Counterdrug Technology Development Program Office has been sponsoring research in facial recognition algorithms since 1993 and has been instrumental in establishing facial recognition capabilities throughout the industry. The FERET program conducted a series of three tests from August 1994 through September 1996, consisting of side-by-side comparisons for different vendors, and test results rated Visionics software as the leading performer in the industry.

Visionics offers these off-the-shelf products:

- FaceIt® DB – Provides the database and search engine components for manual searches within single or multiple databases. Images can be exported from existing databases or can be enrolled (added) using a video camera (video or digital).
- FaceIt Sentinel™ – An on-demand facial recognition program allowing users to either perform a screen grab of a facial image or mouse click on a facial image within a video stream to initiate a search against the FaceIt® database. FaceIt Sentinel™ includes FaceIt® DB.
- FaceIt® Surveillance – Offers all the functionality of FaceIt Sentinel™ but also includes the automatic acquisition of faces from the video stream. FaceIt® Surveillance includes FaceIt® DB.
- Software Developer Kits – A suite of software development tools for integrating FaceIt's capabilities into finished applications and for inserting company-specific customizations into Sentinel™, Surveillance®, and DB®.

The main difference between Sentinel™ and Surveillance® is that once invoked, Surveillance® will continually monitor a video image, seeking to identify faces and compare them against a facial image database. (User preferences control parameters such as: confidence threshold, number of faces to find per picture, search intensity, and redundancy filtering.) With Sentinel™, the user must initiate the find by using the keyboard or mouse. This assessment was conducted using FaceIt Sentinel™ version 2.0.0.12.

Minimum recommended hardware and software requirements for Sentinel™ are:

- Microsoft® Windows NT version 4.0 with Service Pack 5 or higher, or Windows 2000
- 500 MHz Pentium-compatible processor
- 128 MB memory
- 100 MB available hard drive space for installation; additional 400 KB per 1000-face database
- 16-bit color video display adapter
- 800 x 600 screen resolution
- 4x CD-ROM drive
- Video for Windows (VFW)–compatible video capture system with 320 x 240, 16-bit resolution; Peripheral Component Interconnect (PCI) capture system recommended.

The assessment platform consisted of the following configuration:

- Toshiba Tecra 8100 series notebook computer with Microsoft Windows 2000 Professional, Service Pack 1
- 500 MHz Intel mobile Pentium III processor
- 256 MB memory
- 18 GB hard drive
- 14.1" active matrix color LCD display (16,000,000 colors at 1600 x 1200 resolution)
- 24x CD-ROM drive
- 3Com HomeConnect PC Digital Camera, Model No. 3718, with Universal Serial Bus (USB) support.

Although the use of a laptop computer was an implied operational requirement, we discovered that this specific configuration was not ideally suited for the assessment task for several reasons. First, the Toshiba computer uses a USB port as the interface between the digital camera and the computer. This is a relatively new connectivity specification, and although its use is expected to grow, it does not yet enjoy the broad multi-tier support present with older, more established interfaces such as serial, PS/2, or PCMCI ports. Consequently, only two USB digital cameras are endorsed by Visionics. This list is expected to grow with time, but for this assessment, we were limited to one of the two supported home-quality ($150) digital video devices.

Second, Visionics recommends a VFW–compatible PCI video capture system (frame grabber). Frame grabbers are responsible for extracting a single frame image from the continuous video input stream provided by a camera or other video device. Grabbers are typically computer cards and can be installed in any available PCI slot on conventional desktop computers. The Toshiba Tecra 8100 is a notebook computer and, consequently, has neither a built-in frame grabber nor the ability to install one. This caused difficulty during software installation and resulted in the need for assistance from Visionics technical support. It appears to have ultimately reduced the number of configuration options available.

Finally, Visionics' products were initially developed for the Microsoft® Windows NT version 4.0 operating system and have only recently become available to run under Windows 2000. Visionics continues to provide support for NT products while simultaneously expanding and migrating its product line to operate in the Windows 2000 environment. As with the USB port issue mentioned previously, this has been an operation-in-progress and is not yet complete.

Nevertheless, installation has been completed successfully and the product functioned as expected. None of the issues mentioned previously is insurmountable. Individually, they would not have warranted mentioning and are noted here only because together they represented a substantial effort to install and configure the test environment. In the short term, these problems can be avoided simply by selecting a system using older technologies. In the long term, maturation of second and third party hardware and software products will broaden support bases, and issues like these are expected to disappear.

Some of the more significant configuration parameter settings used during product assessment included:

- redundancy filter: off
- auto recognition: on
- express search: on
- express enroll: on
- confidence threshold: 5
- digitizer setting is 352 x 288 x 24-bit resolution.

Out of the box, Visionics does not provide an easy way to invoke Sentinel™. The most direct method requires the following steps:

1. From the desktop, double-click on the application icon to invoke the FaceIt Sentinel™ module.
2. Click on "Surveillance."
3. Click on "Gather Images."
4. From a pull-down menu, select the name of the database against which subsequent searches will be conducted. (Note that Sentinel™ remembers this value, and if it is the same database as specified in the last session, this step can be skipped.)
5. Click "OK."

An evaluation copy of a generalized macro-scripting tool, Macro Express Version 2.1c (Shareware, $30 – http://www.macros.com), was installed. This allowed the user to invoke Sentinel™ much more quickly and easily. Macro Express provides the capability for users to create a desktop icon that can automatically execute all of the steps listed above simply by double-clicking on a single desktop icon. One-click execution of the program is a desirable feature.

Unfortunately, the method by which Macro Express performs its function requires that each computer using the icon purchase a Macro Express license and install it on the resident computer. This method could prove cumbersome in a multi-system remote location deployment scenario, and alternative methods should be investigated. One approach would be to evaluate the Visionics software development kit to determine if such capability can be developed and deployed solely within the Visionics framework.

3. TESTING APPROACH

An empty database was created using FaceIt® DB and accepting default database initialization parameters. Once established, the database was populated with images of local staff. All images were captured in the same physical location under identical conditions. The camera was connected to a tripod, adjusted based upon the height of the subject, and situated 24–30 inches from the subject. Lighting included floor-to-ceiling windows (see the following diagram), a 40-watt fluorescent desk light directly behind the camera, and five 40-watt fluorescent ceiling lights distributed roughly equally about the room. Skies were partly cloudy to sunny. Lighting was not configured according to Visionics' recommended optimal specifications (see Appendix A, *Visionics White Paper: Biometric Facial Recognition Image and Lighting Information*) but does more accurately represent real-world conditions.

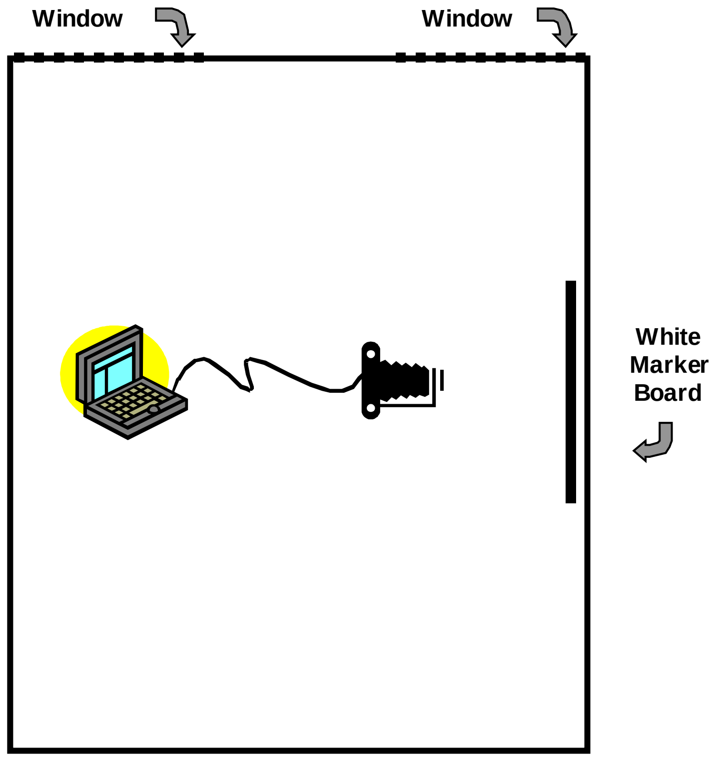

Studio Configuration

The software product used automatically assigns each image a quality score, rating such image elements as darkness and focus. Scores ranged from -1 to 10, with -1 indicating that the test could not be performed and 10 being the best score. The following elements were evaluated:

- Head size — the number of pixels between the eyes. If 25 or more pixels were present, a score of 10 was given.
- Cropping — the amount of face still visible in the cropped canonical image. If the entire face was visible, a score of 10 was given.
- Brightness — the saturation level of an image resulting from brightness. A low score indicated that the image was overexposed because the subject was too close to a light source, or because of the strength of the light source. A low score can also indicate a reflection on the person's forehead.
- Darkness — the saturation level of an image resulting from darkness. A low score indicated the image was underexposed because the subject was too far from the light source, or because of the strength of the light source. Brightness and darkness scores are the most critical for face recognition performance.
- Focus — the contrast of the image. This element value was weighted more toward the middle of the face. If an image had high contrast, a score of 10 was given. A lower score may indicate that the camera was not in focus.

- Eyes — the visibility of both eyes. If a subject was not wearing glasses and both eyes were visible, a score of 10 was given. If the eyes were not found, a score of -1 was given. A score between -1 and 10 indicated the eyes were partially obstructed, possibly because of glasses or the position of the head.
- Glare — the level of glare on the glasses. If the "eyes" score was less than 10, it was assumed that the person was wearing glasses. If the eyes were not found, a score of -1 was given.
- Overall quality — the lowest positive number of the above scores. Overall quality was used to determine the level of image quality:

Overall Quality Score	**Image Quality**
0 – 2.5	Poor
2.5 – 5	Fair
5 – 7.5	Moderate
7.5 – 9.5	Good
9.5 – 10	Excellent

Of the 40 images gathered, 7 were ranked as having excellent image quality, 26 had good image quality, 6 had moderate image quality, and 1 had fair image quality. There were no poor-quality images. Moderate quality images had lower scores on glare and focus, with difficulty split equally between either glare or focus.

Twenty-nine out of 40 (73%) of the subjects were male and 20 out of 40 (50%) of the subjects wore glasses. One subject wore a baseball cap and was enrolled twice – once wearing the cap and once without. Eight men had mustaches and three had partial or full beards.

The technology assessed uses the centers of the eyes as indexing and referencing points. When a head-like object is visible, the software uses a variety of pattern-matching algorithms to determine if a face is indeed present at that location. These algorithms are capable of accurately detecting the simultaneous presence of multiple faces and determining their precise positions.

After a face has been detected, the system locates the eyes within the head. The centers of the eyes are used as alignment points for locating and indexing all other references. This product can be configured to align with the eyes automatically, or individual manual adjustments can be made to improve accuracy. Automatic alignment was used for this assessment, and although results were quite good, the system did have difficulty locating eyes when the subject did not look directly into the camera, whenever lighting was inadequate, or if the subject wore glasses.

The facial image is extracted from the background and subjected to a number of proprietary preprocessing stages to compensate for size, lighting, expression, and pose. The "normalized" face is subsequently transformed into an internal representation called a faceprint using a mathematical technique called Local Feature Analysis. This digital code contains information intrinsic to the face, and it is used to determine facial identity by matching the live faceprint against a database of stored prints of known individuals.

In surveillance mode, the system extracts an image from the surveillance device, develops a faceprint, and searches the database for entries most resembling the image being searched. Each image is ranked on a scale of 1 to 10. Thumbprint representations of the most likely candidates are returned in order of relevancy, and the user can then compare each individually to determine if a match has been made. Extensive customization allows users to control such values as:

- maximum number of thumbprints returned

- minimum relevancy threshold for thumbprints returned – based upon the ranking scale.
- face-finding parameters (head size, low resolution, confidence threshold, etc.)
- friend/foe/any audio alarms
- redundancy filter – to prevent the system from attempting to search on the same individual over and over.
- auto recognition – determines if the database is automatically searched every time a new image is captured.

Images captured during surveillance can be automatically enrolled to the image database or can be enrolled individually (manually). Images can be enrolled in the database being searched or into a separate database.

4. RESULTS

After the initial database was built, the 3Com digital video camera was used to capture images that were subsequently searched for in the database. Performance metrics measured included the system's ability to match faces under similar and/or varying conditions and the confidence level with which matches were made.

In the first (control) series of tests, images were taken in the same environment in which the original database images were captured. This series included individuals whose images were present in the database as well as those who were not. Variations such as adding/removing glasses and/or caps were also included (see following table).

Control Environment

Subject	Status	Rank	Confidence
1 (George S.)	present/no change	1	9.2
2 (Steve F.)	present/no change	1	7.9
3 (Steve P.)	present/glasses removed	3	6.5
4 (Jim T.)	present/glasses removed	1	7.7
5 (Joe T.)	present/glasses added	1	8.6
6 (Jackie W.)	present/glasses added	1	6.8
7 (Rich B.)	present/cap added	no face found	N/A
8 (Mike M.)	present/cap added	1	7.2
9 (Ben T.)	not present	N/A	#1 was 7.3
10 (Patsy P.)	not present	N/A	#1 was 7.5

In all but two instances, the subject was identified. In one case, the subject's glasses were removed, and although the correct image was returned as one of the thumbnails, it was ranked third in confidence. Repeated tests with this subject yielded similar results.

In the other instance, the subject was wearing a cap, and the product was unable to even find a face to index. It is assumed that the difficulty was a lighting problem because the bill of the cap was shading the subject's eyes.

It is also notable that tests involving subjects who were not in the database (subjects 9 and 10) returned rankings higher than some tests in which the subject was found. This suggests that system performance may appear artificially high because of the relatively small number of images in the database. Although the subject was identified, confidence rankings were lower than desirable, and if the database had contained a large number of subject images, a mismatch may have resulted.

The second (field) series was performed in a separate room with no windows and only overhead lighting present. System hardware and software, as well as the subjects and their statuses, were identical to those in the first series. Besides lighting, the other major difference between the two environments was that while the picture background of the control environment was a simple white marker board, the subject background for the field environment was a technology poster. The poster's background was dark blue (almost black) and populated with multiple text panels and computer-screen representations of various sizes. Most text panels were yellow with medium blue text, but the computer screen representations included red and green segments as well.

The system had significantly more trouble in this environment than in the first. In all cases except one, results were worse. The single exception was subject 3. Although this was the only instance in the control series where the number one rank returned was not the original subject, in the field series the correct subject was returned and with a higher (7.9) rank than in the control set (6.5). Equally surprising were subjects 2 and 4 because the control set returned relatively high confidence levels on each (7.9 and 7.7, respectively), and in the field series, the original subjects were not even ranked. Finally, subject 9 was surprising in that the product could not even find a face within the image. Subsequent ad hoc testing revealed that subject 9's face could be found if the subject first removed his eyeglasses.

Field Environment

Subject	Status	Rank	Confidence
1 (George S.)	present/no change	1	7.2
2 (Steve F.)	present/no change	did not return	N/A
3 (Steve P.)	present/glasses removed	1	7.9
4 (Jim T.)	present/glasses removed	did not return	N/A
5 (Joe T.)	present/glasses added	1	7.3
6 (Jackie W.)	present/glasses added	1	6.5
7 (Rich B.)	present/cap added	no face found	N/A
8 (Mike M.)	present/cap added	1	6.7
9 (Ben T.)	not present	no face found	N/A
10 (Patsy P.)	not present	N/A	#1 was 6.7

In general, results were somewhat disappointing. Although the product did an impressively good job of managing such abstract objects as facial images, it remains to be seen whether the technology is sufficiently advanced to be useful in practice. The main factor in determining the system's ability to recognize faces seems to hinge upon factors external to the software. For example, a database populated with excellent-quality images and a video monitoring system, which can capture high-quality face-forms, should have a much higher success rate than one that is searching against only mediocre-quality images or is trying to match a poor-quality face capture to a high-quality database.

Image quality is highly dependent upon lighting and distancing conditions, camera quality, computer hardware configurations, and pose (the subject should be looking directly at the camera when the image is captured) — all or none of which may be controllable by the user. In situations where images are captured in a preconfigured, prepared setting, image quality can be good. On the other hand, portable systems may not have the luxury or ability to control these factors and may not even have the ability to obtain a good subject pose. Any consideration of implementing a facial recognition system must take these items into consideration in the context of how the system ultimately will be used.

5. CONCLUSIONS

The product's ability to accurately match and retrieve images is dependent upon the software's accuracy in indexing and cataloging both the images stored in the database as well as those captured during surveillance. Consequently, camera quality and lighting conditions are probably the major factors in determining product usability. As mentioned previously, the hardware/software configuration used here did not permit broad experimentation with different camera types. This was seen as a significant constraint in our assessment. Lighting also was a concern, and although this study did not gather images under optimal lighting conditions, it did take place in a controlled environment that may not be feasible in a remote operation deployment scenario.

Camera utility also is associated with camera quality. The 3Com unit was designed as a home desktop unit and did not lend itself well to portability. The camera was awkward to hold and transport and easily lost focus if moved. Measuring about 4" x 2.5" x 1.5" and having a 9-ft attachment cable, it certainly failed to meet the criterion for unobtrusiveness. Any realistic implementation must take into consideration the latest developments in miniature digital camera technologies and how well they can be integrated into new/developing facial recognition applications.

Constructing implementation scenarios that involve deployment of multiple copies of a centrally located database is a complex process, especially when the deployed systems also are capturing images that must be integrated back into the central database (and subsequently disseminated back to the remote systems). Consideration must be given to how these processes will be performed. Typical issues might be:

- How often will remote databases be updated? How will the updates take place?
- Will the database be static or will new images be added? If new images are being added, will they be captured by the remote systems, from a separate source, or a combination of the two?
- How will conflicts be resolved? For example, if two remote systems upload images of the same individual, how can the duplication be identified and how will the correct image be chosen?
- If a match is found by one of the remote systems, what, if anything, should be done regarding the central system and/or the other remote systems?
- What are the security implications?

These issues can be particularly difficult and should be given extensive consideration well before implementation actually begins.

6. RECOMMENDATIONS

If interest warrants further investigation, it is suggested that the next step be to analyze to determine and document the specific conditions under which a production system might be used. At a minimum, the analysis should:

- describe in detail and/or provide examples of images that will comprise the database and, if applicable, the methods by which new images will be added;
- describe and/or make available an environment with layout, furnishings, concealment requirements, lighting, and subjects representative of that which will be used by remote systems; and
- determine threshold performance and confidence limits regarding acceptable hit/miss ratios.

The information obtained from the analysis can then be used to obtain one or more systems closer to that which will be used in the operationally functional state. A more realistic evaluation should then be made by using information gained during this assessment to build a larger database, consisting of images captured and managed as they would be in a real-world situation.

Appendix A

Visionics White Paper:
Biometric Facial Recognition Image and Lighting Information

Biometric Facial Recognition Image and Lighting Information

The following document contains information pertinent to the ***optimal*** image and lighting conditions required for facial recognition. These are the parameters that should be used in circumstances where the capture and creation of facial images can be performed in a controlled environment with subject participation.

Images created using these parameters should perform better in both one-to-many (identification) and one-to-one (verification) matching scenarios.

Image Quality Specifications

The purpose of this description is to define "Driver's License" quality images for Visionics face recognition and face-finding technology. Deviations from high quality images may affect the performance of the technology and should, therefore, be avoided if possible.

The required image format is JPEG, with 24-bit color or 256 shades of gray color depths. The JPEG compression for the photos should be, at maximum, roughly 15:1 ("quality" 80 for JPEG).

The best images for processing are 300 x 400 x 24-bit resolution. The entire head, including hair and part of the neck, should constitute the bulk area of the image, and there should be roughly 100 pixels of data from eye to eye. (The minimum quality image we recommend is 100 x 125 x 256 Grayscale JPEG with 15:1 compression, and 25 pixels from eye to eye.)

The best faces for matching are directly facing forward under controlled, balanced lighting, with eyes open and a neutral expression like a mug shot photograph.

Under no conditions should the head/neck be cut off in any way, as this will severely compromise the automatic face finder.

Other situations to be avoided include glare on eyeglasses that obscure the eyes, sunglasses, closed eyes, mouths open during speech, strong smiles with exposed teeth, and variations in pose (left/right) or tilt (up/down) beyond 10 degrees in any direction.

Larger resolution photos can be accepted but generally do not increase accuracy to any measurable extent and take longer to process. Smaller photos are also accepted but accuracy decreases dramatically for sizes below 200 x 300. The absolute MINIMUM quality would be 100 x 125 x 256 Grayscale JPEG with 15:1 compression, and 50 pixels from eye to eye.

Lighting Configuration

The goal of the lighting specification is to create a uniform wall of lighting, which should reduce lighting-dependent effects upon the face. Also, lighting should come primarily from the front of the subject so that shadows from overhead lighting do not confuse the facial enrollment algorithm.

It may be possible to duplicate the optimal lighting conditions using photoflash equipment rather than fixed photo lighting. This process has not been thoroughly tested by Visionics and, therefore, the following is based on the fixed-lighting scenario.

Two halogen lamps with umbrella directors should be used to provide lighting. A light meter can be used to ensure that acceptable lighting is directed upon the subject. When measuring the lighting, the light meter sensor should face the camera, parallel to the backdrop and perpendicular to the floor. Measure the light levels in the area in which faces may be found — 12 inches above and below the face position for the average person. This will ensure that light levels will be acceptable for all subjects. The level of light should be 700 lux, +/- 20 lux.

Position the light meter at the top of the head height range and direct the sensor toward the ceiling. The sensor should now be parallel to the ground. Use the light sensor to measure the overhead lighting. The level of overhead lighting must be less than the lighting coming from the front lights.

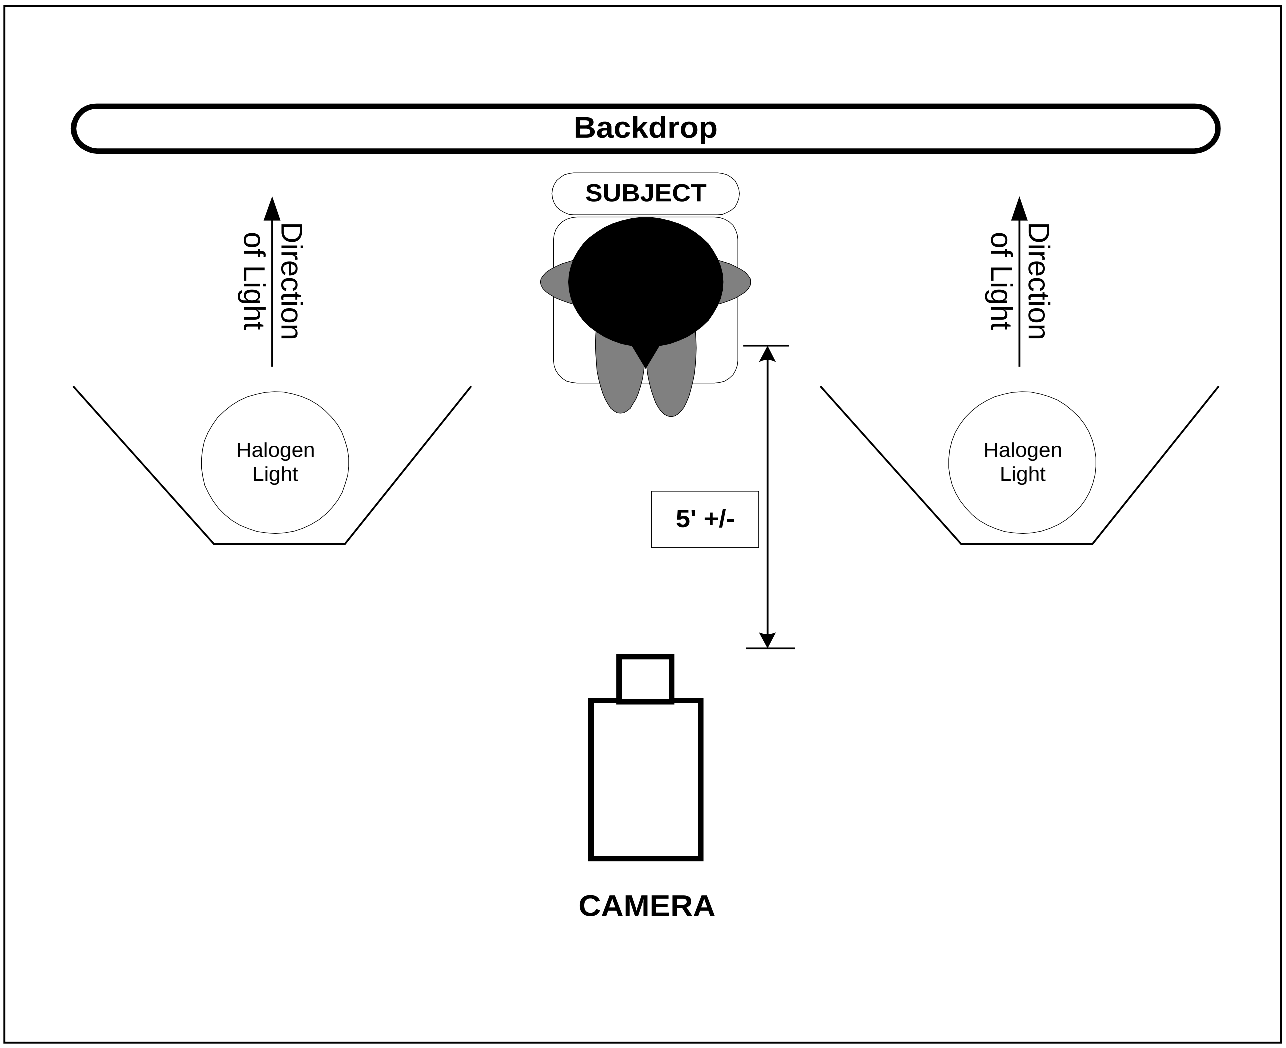

The Image Quality module evaluates the quality of an image based on the following criteria:

Face Finding. This indicates the general quality of the face found in an image without concern for the remaining quality issues.

Head Size. This indicates whether the subject is too far away from the camera. The operator can reposition the subject or adjust the camera lens to compensate.

Cropping. This indicates whether the subject's entire head (face) is in the frame. This can identify problems that are resultant of a taller or shorter subject where part of the subject's face is cropped because the subject is too close to the camera. The operator can tilt and/or pan the camera to compensate.

Note that it is possible for the subject's head to be both too small and simultaneously cropped. For example, the subject could be standing far from the camera and to the side.

Brightness. This indicates that the subject is overexposed, possibly because he or she is standing too close to the light source. On the other hand, it may be a result of the light source's strength. The operator can reposition the subject or adjust the light source to compensate.

Note that it is possible for a subject to be both overexposed and underexposed simultaneously in different parts of the face. For example, a strong light source directly overhead may cause the forehead to be very brightly lit and simultaneously create dark shadows over the eyes.

Darkness. This indicates that the subject is underexposed, possibly because he or she is standing too far from the light source. On the other hand, it may be a result of the light source's strength. The operator can reposition the subject or adjust the light source to compensate.

Blur. This indicates the degree to which the image is blurry. Blurriness may be caused by the camera being out of focus. The operator can manually adjust the focus on the camera or ensure that auto-focus, if available on the camera, is turned on. For auto-focus use, most cameras will focus on what is in the middle of the frame. If the subject is off to the side, auto-focus may not work.

Eyes Clear. This indicates whether the subject's eyes are obstructed, potentially because the subject is wearing glasses. For the purpose of face recognition, it is better if the subject does not wear glasses. However, your implementation may make it impossible or impractical to remove the subject's glasses. If this is the case, you should ignore this value.

Note that this is essentially a Boolean test, so the value will be –1, 0, or 10. A value of 10 indicates that there were no eye obstructions found, meaning this is a higher-quality image for face recognition. Zero indicates that there was an eye obstruction, potentially glasses, and –1 indicates that the image is not suitable for this test. This test is currently only valid on high-resolution images. So, this test will be performed if the result you get from the **Head Size** test is less than approximately 8.5.

Glare Free. This indicates glare — generally from a subject wearing glasses — obstructing the subject's eyes. The operator can reposition the subject or adjust the light source to compensate.

Note that this is essentially a Boolean test, so the value will be –1, 0, or 10. A value of 10 indicates that there was no glare found, meaning this is a higher quality image for face recognition. Zero indicates that there was glare found in the image, and –1 indicates that the image is not suitable for this test. Test is only valid on images with eye obstructions. If the person passes the **Eyes Clear** test, then this test will not be performed.

Quality Confidence. This is the overall quality of the image. It is simply the lowest of all scores for all but the last two (**"Eyes"** and **"Glare"**) quality tests that are used. You or your client application most likely will examine this value first and then decide whether to pay attention to the individual scores.

Video Input FAQs

Question	Answer
2.1 What video device standards do you support on the Windows platform?	*We currently support in our products the Microsoft Video For Windows (VFW) standard. In the near future (Q3 99) we will also support the Direct Show standard. Nearly all video input devices designed for the windows platform support VFW.* *[Using our low-level API calls provided by our SDKs however, a programmer can connect any video device to our algorithms.]*
2.2 Do you support USB cameras?	*If the camera manufacturer provides a VFW software driver for its device, then that device will work with our products and developer software.* ***Our software does not connect directly to hardware.***
2.3 What capture hardware and cameras do you recommend for use with your technology?	***PCI capture cards*** *The Integral Technologies FlashBus MVLite PCI capture card:* www.integraltech.com *The Winnov PCI capture card:* www.winnov.com *works reasonably well with the Winnov camera and takes standard video input as well.* *The 3COM big picture PCI card and camera is a great low-cost system:* www.3com.com ***USB Desktop Cameras*** *The Kodak DVC323,* www.kodak.com *The Winnov USB camera:* www.winnov.com ***Parallel Port Cameras*** *The Vicam camera:* www.vistaimaging.com ***Desktop Cameras with video output*** *For a large list, see the Netmeeting web site:* www.microsoft.com*/windows/netmeeting* *The Toshiba CCD desktop camera is a good example of a high-quality/low-cost desktop camera:* www.toshiba.com *We also like the Howard HA6800 with Zoom lens for high-quality video input for desktop systems: see* www.visionics.com *For demos, teleconferencing, or high-end use, we like the more expensive Sony EVI-D30 pan tilt zoom auto-focus camera:* www.sony.com *The Connectix Quickcam series of cameras is not recommended for use with face recognition due to below-average performance.* ***Fixed-field Surveillance Cameras*** *The Philips LTC 450 with Philips lenses provides excellent output:* www.philipscss.com. *For demos or indoor fixed-field surveillance usage, we also like the Sony EVI-D30 pan tilt zoom auto-focus camera:* www.sony.com ***Pan-tilt-zoom Surveillance Cameras*** *Application Specific*
2.4 What sort of image enhancements might be	*You should enable automatic gain control on the video camera.*

Question	Answer
required based on the video input?	
2.5 How can the quality of video be controlled to ensure optimal results?	*Avoid including a bright light source in the video field of view, such as the sun, or, when indoors, a window in the background field of view.* *In general, avoid situations that will generate a photographic backlighting problem.*
2.6 What are the recommended video digitizer settings?	*Analog video input must be digitized (resolved into pixels) before it can be processed for face finding and face recognition. The default digitizer setting for most desktop systems is 160x120 pixels.* *This resolution is not recommended for use with Visionics technology. Instead we recommend:* *For desktop verification:* ***RGB888 320x240 pixels.*** *For video surveillance:* ***RGB888 640x480 pixels****.* *The setting can be adjusted via the manufacturer's digitizer options dialog box, which is a standard component of VFW software drivers.*
2.7 What is the full definition of best quality video input for face finding and face recognition?	*The video digitizer settings are RGB888 640x480.* *The face and neck are clearly visible in the video field, and roughly the face takes up one-third of the video field. Under no conditions should the head be cut off in any way, as this will compromise the abilities of the automatic face finder.* *The face should not be so close to the camera that it causes a "fish-eye" effect where the nose is larger than normal.* *The best faces for matching are directly facing forward under controlled, balanced lighting, with eyes open and a neutral expression, like in a mug shot photograph.* *Other situations to be avoided include glare on eyeglasses that obscure the eyes, sunglasses, closed eyes, mouths open during speech, strong smiles with exposed teeth, and variations in pose (left/right) or tilt (up/down) beyond 15 degrees in any direction.*
2.8 What is the minimum video input specifications required for the Visionics technology to maintain effectiveness?	*There should be a minimum of 20-25 pixels between the eyes. The person's pose should be within 30 degrees of frontal, within 15 if possible.*

DISTRIBUTION

1. Department of Energy – IN
2–6. Department of Defense
7. Ron Cain, 1099 Commerce Park, MS-7620
8-10. George Singleton, 1099 Commerce Park, MS-8260
11. Advanced Computing Technologies Resource Center, 1099 Commerce Park