IMTC 2005 – Instrumentation and Measurement
Technology Conference
Ottawa, Canada, 17-19 May 2005

Effect of Image Compression on Iris Recognition

Robert W. Ives, Bradford L. Bonney, Delores M. Etter
EE Department, U.S. Naval Academy, Annapolis, MD 21402

Abstract *– Iris recognition is a proven, accurate means to identify people. Commercial iris recognition systems are currently employed to allow passengers in some airports to be rapidly processed through security, to allow access to secure areas, and for secure access to computer networks. With the growing employment of iris recognition systems and associated research to support this, the need for large databases of iris images is growing. If required storage space is not adequate for these images, compression is an alternative. It allows a reduction in the space needed to store these iris images, although it may be at a cost in some amount of information lost in the process. This paper investigates the effects of image compression on iris recognition. Compression is performed using JPEG2000, and the iris recognition algorithm used is based on several methods, including the Daugman algorithm.*

Keywords *– Iris recognition, Daugman, Hamming distance, JPEG2000.*

I. INTRODUCTION

Biometric identification or verification of identity is currently a very active field of research. Many applications that require some degree of confidence concerning the personal identification of the people involved, such as banking, computer network access or physical access to a secure facility, are moving away from the use of paper or plastic identity cards, or alpha-numeric passwords. These systems are too easy to defeat. A higher degree of confidence can be achieved by using unique physical and/or behavioral characteristics to identify a person; this is biometrics.

In order to use biometrics for identification, the biometric data must be collected by some means from the individuals in question. In some cases, this may be a costly and time-consuming process, and the data obtained is valuable and must be protected. Additionally, data collections can create an inordinate amount of data that puts a strain on the available storage. To alleviate this problem, one available option is compression. In many applications where compression is required, but no loss of information is acceptable (such as monetary transactions or some medical applications), lossless compression is necessary; that is, compression without loss of information.

There are many lossless compression algorithms available that work best on certain types of data, such as predictive coding for one-dimensional waveform data and string coding for text. For imagery, JPEG2000 and lossless-JPEG have demonstrated very good lossless compression performance with most types of imagery. Unfortunately, lossless compression has a major drawback in that the reduction in file size is on the order of only 1.5:1 to 3:1 for many types of imagery. On the other hand, these algorithms can readily compress data further if some loss of information is tolerable. It is up to the user of the data to determine how much loss of information is acceptable.

The iris (see Fig. 1) is the colored portion of the eye that surrounds the pupil. Its combination of pits, striations, filaments, rings, dark spots and freckles that is evident under near-infrared (NIR) light make for a very accurate means of biometric identification [1]. Its uniqueness is such that even the left and right eye of the same individual is very different.

In this paper, we investigate the effects of lossy compression on the ability of an iris recognition system to accurately identify individuals. The performance is evaluated by means of the change in Hamming distances between IrisCodes using an iris recognition implementation based on several algorithms, including the Daugman algorithm [1]. Typically, a database for an iris recognition system does not contain actual iris images, but rather it stores a binary file that represents each processed iris, such as Daugman's IrisCode, stored as 512 bytes per eye. The size of such a database may not necessarily be prohibitive. However, we do not propose compressing this template data, but instead the original images from which they were created. We seek to compress the original imagery because it is this data that is valuable, and serves as training and testing imagery for the development of new algorithms. Its importance became apparent to the authors as we began to collect our own iris database, which is discussed in the next section.

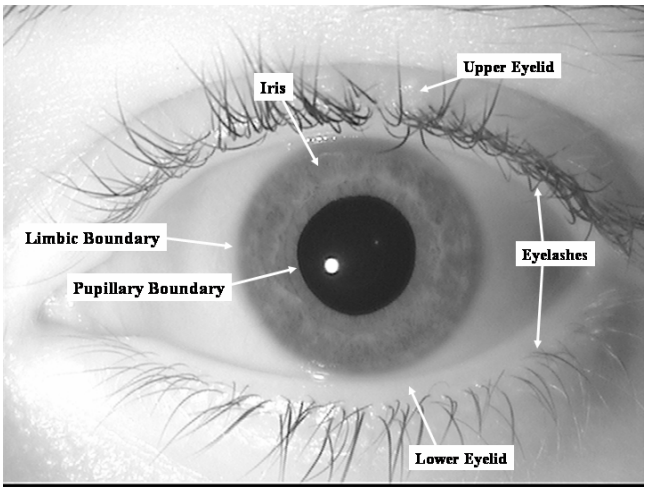

Figure 1: An Example Near-Infrared Iris Image

Compression has been investigated and used in some biometric applications, such as the FBI standard for fingerprint compression [2]-[3], or using MPEG compression [4]-[5] for video that may be used in facial recognition applications. There has been some limited research in the area of iris image compression [6], but this was compression applied to IrisCodes, not iris images. Here, we address the issue of compression applied to the iris imagery itself.

II. JPEG2000

JPEG2000 is the new compression standard published by the Joint Photographic Experts Group [7]. It employs state-of-the art compression techniques based on wavelet technology. Like the previous JPEG standard, it allows for both lossless and lossy compression of imagery. Lossy compression means that some information is lost in the process, and the amount of information lost is dependent on the algorithm used for compression, as well as the amount of compression desired (that is, the size of the compressed file).

JPEG2000 offers some advanced features, such as region-of-interest (ROI) coding, where the user could identify regions of the image that should be compressed to a higher quality than the surroundings. ROI coding might prove advantageous in iris image compression, since it would allow the iris itself to be compressed with less loss of information than other areas of the image that are not used in recognition. For this research, both lossless and lossy compression of iris images were tested using the default parameters and options for JPEG2000. JPEG2000 was implemented using Win32 executable code freely available from Kakadu Software [8].

Fig. 2 displays an original iris image before and after compression to 20:1 using JPEG-2000. The original image was collected with the LG IrisAccess 3000 system. Comparing the original and the compressed image closely will reveal some detectable differences, primarily in the areas of high detail in the original image where compression artifacts or smoothing is noted. Statistically, the two images are not very different; the maximum difference between the two images is 26 gray levels, and the overall average difference is 0.056328 with a standard deviation of 2.951321. Overall, JPEG-2000 does a good job of maintaining the detail information even up to a compression of 20:1.

III. IRIS RECOGNITION ALGORITHM

Commercial iris recognition systems today use the algorithm developed by John Daugman [1]. This patented algorithm is not available for free use, so an alternative for research purposes can be found in the implementation created by Libor Masek [9]. This algorithm follows the Daugman algorithm to some extent, but also incorporates parts of other reported algorithms. Most notably, the MATLAB code is freely available [10].

The Masek algorithm differs from the Daugman algorithm in several areas. This includes the use of the Hough transform

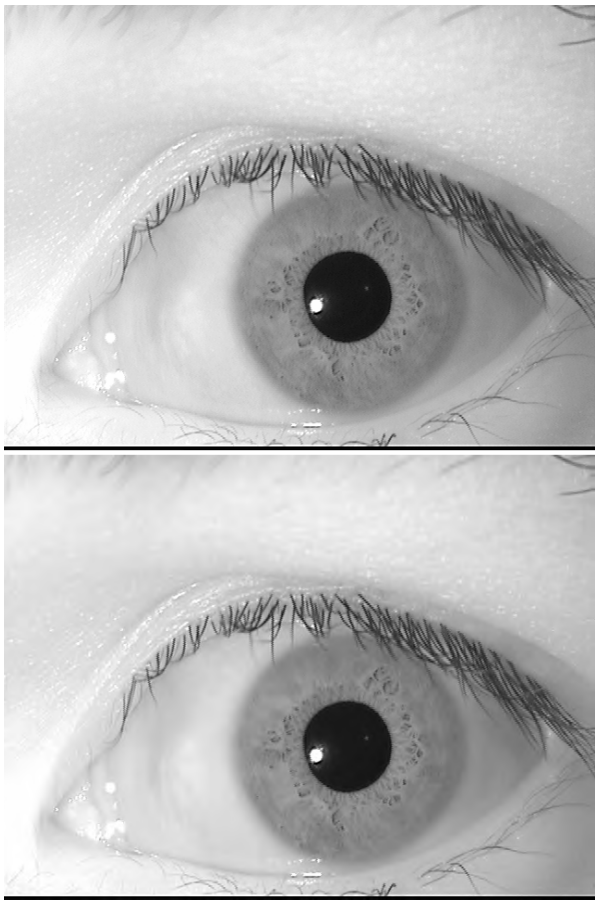

Figure 2: An original iris image (top) and compressed to 20:1 with JPEG-2000 (bottom).

to detect the circular inner iris boundary (the pupil) and outer iris boundary and its use of Log-Gabor wavelets vice Gabor wavelets for feature coding. When an image is input to the algorithm, the output is composed of two parts: the phase-code bits that represent the distinct patterns within the iris; and a mask which represents the locations of iris pattern bits which are used to compare irises, as opposed to noise that is present in the image among the iris patterns (such as eyelashes, glare, etc.), but should not be used for comparison.

IV. METHODOLOGY

The images used in this research come from two sources. First, we used the Chinese Academy of Sciences (CASIA) iris database [11]. This is composed of images of 108 different eyes, with 7 images of each eye (totaling 756 iris images). These images are 320x280 8-bit bitmapped images (.bmp), each occupying 92,160 bytes on a hard drive. A second database was comprised of images collected using the video output of our lab's LG IrisAccess 3000. This video was fed to a Matrox Meteor II frame grabber installed in a Dell Dimension 4600 desktop computer. Using the MATLAB Image Acquisition toolbox, the video was piped directly into

a graphical user interface (GUI) that runs on MATLAB 7.0. Using this GUI, the user is set to capture video at 10 frames/sec for one second (these numbers are based on our desire to capture 10 images of the same iris). Each frame becomes a 640x480 8-bit bitmapped image (.bmp), and occupies 309,248 bytes on the hard drive.

Performance was measured by observing the effect on fractional Hamming distances between the IrisCodes from the original and decompressed images, computed using the Masek algorithm [9]. The fractional Hamming distance (HD) between IrisCodes A and B is defined as:

$$HD = \frac{\left\| (\text{code A} \otimes \text{code B}) \cap \text{mask A} \cap \text{mask B} \right\|}{\left\| \text{mask A} \cap \text{mask B} \right\|}. \qquad (1)$$

The $\otimes$ operator is the Boolean XOR operation to detect disagreement between the pairs of phase code bits in the two IrisCodes (code A and code B), and mask A and B identify the values in each IrisCode that are not corrupted by artifacts such as eyelids/eyelashes and specularities. The $\cap$ operator is the Boolean AND operator. The $\| \cdot \|$ operator is used to sum the number of "1" bits within its argument. The denominator of (1) ensures that only the phase-code bits that matter are included in the calculation, after any artifacts are discounted. This serves as a measure of recognition performance, as it is the fractional Hamming distance that determines if identification has been made. A value of HD = 0 indicates a perfect match between the IrisCodes, while typically a Hamming distance of ≤ 0.32 allows identification with high confidence and is here used as a threshold for recognition.

We compressed 44 images from each database using lossy JPEG-2000 to compression ratios of 4:1, 6:1, 8:1, 10:1, then 11:1, 12:1, etc. up to 20:1 and their IrisCodes were created. As mentioned in Section II, when using JPEG-2000, the default compression parameters were selected; the only option chosen was the desired bit rate (i.e., compression ratio). For each original iris image, there were 14 compressed versions, which populated each database with 660 images (44 x 15). To derive the performance results, each original iris image was compared against every other image in the database. This means that a total of 28,996 comparisons were made (659 x 44), of which 616 comparisons (44 x 14) were enrollee attempts (the irises should match) and 28,380 (44 x 43 x 15) were imposter attempts (the irises should not match).

V. RESULTS

To form a baseline regarding compression of iris images, JPEG-2000 was used first to compress the iris images without loss of information. Lossless compression allows exact reproduction of the original image from the compressed file. Depending on the algorithm used, the size of the compressed file will vary. In addition, different images will result in different compression attainable when using the same algorithm. The lossless compression results are summarized in Table 1 for each database. Here, the average lossless compression ratio achieved for the 44 images of each database are presented.

Table 1: Lossless Compression by Database

	CASIA	LGIris
Average Lossless CR Achieved	1.74	2.188

Lossy compression effects on recognition performance were evaluated using the False Acceptance Rate (FAR) and False Rejection Rate (FRR), with results summarized in Table 2. A value of HD $\leq$ 0.32 was used to determine whether a match had been made. In computing these values for the LGIris database, there were zero matches made out of 28,380 imposter attempts (images that should not have matched), but four false rejections out of the 616 enrollee attempts (images that should have matched). For the CASIA database, there were 38 false rejections and 0 false matches using the same number of attempts as for the LGIris database. The definition used in Table 2 for the FRR is defined in [12] as

$$(\%)FRR = \frac{\text{\# of incidents of false rejections}}{\text{total \# of samples}} \times 100\% \qquad (2)$$

and the FAR is defined as

$$(\%)FAR = \frac{\text{\# of incidents of false acceptance}}{\text{total \# of samples}} \times 100\%. \qquad (3)$$

In these formulas, for each database, the denominator is 28,996, as stated in Section IV.

Table 2: Recognition Performance Results by Database

	CASIA	LGIris
FAR (%)	0	0
FRR (%)	0.131	0.0138

Concerning the false rejection rate, it is important to note that three of the four false rejections in the LGIris database were associated with the same iris image, while in the CASIA database, three of the images resulted in 33 of the 38 false rejects. One of the CASIA images (image number 064_1_1) resulted in 14 false rejections in 14 attempts. This image is shown in Fig. 3.

Closer inspection of this image reveals some degree of blurring of the iris as well as capture artifacts (noticeable on the eyelashes), not to mention the occlusion of the iris by the upper eyelid and eyelashes. We attribute the poor results using this eye to the image quality. Overall, the quality of the LGIris database imagery is superior to the CASIA imagery.

Typical HD results using the LGIris database are illustrated in Table 3, here for an iris image labeled "Iris00001." The left column denotes the compression ratio

applied to test images to which the original Iris00001 is compared. The middle column displays the Hamming distance computed when the IrisCode for the original Iris00001 was compared against itself and also against compressed versions of itself. The right column was derived by comparing the original Iris00001 IrisCode with an uncompressed image of a different eye (referred to as "Iris00002"), as well as compressed versions of Iris00002.

Table 3: LGIris Database Hamming Distances (HD)

Compression Ratio	Iris 00001 (same eye)	Iris 00002 (different eye)
None	0	0.47078
4:1	0.04188	0.46853
6:1	0.14671	0.46335
8:1	0.12487	0.46339
10:1	0.12420	0.46480
11:1	0.08567	0.47090
12:1	0.08791	0.46882
13:1	0.12366	0.46975
14:1	0.12961	0.46978
15:1	0.13414	0.47019
16:1	0.10559	0.46953
17:1	0.11156	0.46841
18:1	0.11409	0.46666
19:1	0.15485	0.46533
20:1	0.19828	0.46798
Note: HD ≤ 0.32 determines recognition		

VI. CONCLUSIONS

From these results, JPEG-2000 has proven to be a very capable lossy compressor of NIR iris imagery. There was no effect on the false acceptance rate, and only a very slight effect on the false rejection rate. This is noteworthy, given the relatively high compression ratios these images were subjected to. Overall, this means that iris database storage could be reduced in size, possibly by a factor of 20 or even higher (since 20:1 was the limit of compression in this study), and have only a very minor affect on system performance. Further analysis of the false rejections is warranted, and research into how these results scale to a larger database is in progress.

As a state-of-the-art lossless compressor, compression of these iris images using lossless JPEG-2000 could reduce the required storage for a database to approximately ½ of its original size. This may be sufficient in some cases, but significant improvement can be achieved with lossy compression.

Further testing using JPEG-2000 is feasible and in progress to determine additional limits. One feature of JPEG-2000 that was not incorporated in this research was the use of regions of interest. A priori knowledge of a region of interest that should be preserved with less information loss should improve these results. For example, determination of the pupil's location, a relatively simple task in iris preprocessing, would allow identification of an area of interest such that the eye portion of the eye image could be preserved with better quality than surrounding areas (such as eyelids, forehead, etc.). In addition, other options of JPEG-2000, such as choice of wavelet filters can also be examined.

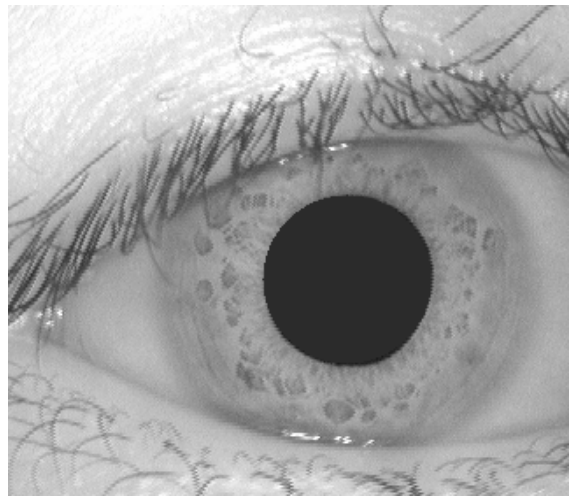

Figure 3: CASIA Image 064_1_1. This image resulted in 14 out of 14 false rejections when compressed.

ACKNOWLEDGEMENTS

Portions of this research used the CASIA iris image database [11] collected by Institute of Automation, Chinese Academy of Sciences. This work has been sponsored in part by the U.S. National Security Agency.

REFERENCES

[1] J. Daugman "How iris recognition works," *IEEE Trans. on Circuits and Systems for Video Technology.*, Vol. 14, No. 1, pp. 21-30.

[2] C. M. Brislawn, "The FBI Fingerprint Image Compression Specification," in *Wavelet Image and Video Compression*, P. N. Topiwala, Ed. Boston, MA: Kluwer, 1998, ch. 16, pp. 271–288, invited book chapter.

[3] J. N. Bradley and C. M. Brislawn, "Compression of fingerprint data using the wavelet vector quantization image compression algorithm," Los Alamos Nat'l. Lab, Tech. Report LA-UR-92-1507, Apr. 1992, FBI report.

[4] J. Daugman, "Face and gesture recognition: overview," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 19, No. 7, pp. 675-676, 1997.

[5] H. Wang and S.F. Chang, "A highly efficient system for automatic face region detection in MPEG video," *IEEE Transactions on Circuits and Systems for Video Technology*, Vol. 7, No. 4, pp. 615-628, 1997.

[6] U. von Seelen, "IrisCode template compression and its effects on authentication performance," Presentation at the Biometrics Consortium Conference 2004, http://www.biometrics.org/bc2004/CD/PDF_PROCEEDINGS/Microsoft%20PowerPoint%20-%20Von SeelenBrief_Update.ppt%20%5BRead-Only%5D.pdf.

[7] The JPEG2000 Standard, http://www.jpeg.org/jpeg2000/index.html, Feb. 27, 2005.

[8] Kakadu Software, http://www.kakadusoftware.com/, Feb. 27, 2005.

[9] L. Masek, "Recognition of Human Iris Patterns for Biometric Identification," M. Thesis, The University of Western Australia, 2003, www.csse.uwa.edu.au/~pk/studentprojects/libor/LiborMasekThesis.pdf, Mar. 26, 2005.

[10] Libor Masek, Peter Kovesi. *MATLAB Source Code for a Biometric Identification System Based on Iris Patterns.* The School of Computer Science and Software Engineering, The University of Western Australia. 2003.

[11] CASIA Iris Image Database, http://www.sinobiometrics.com.

[12] John Woodward, Nicholas Orlans and Peter Higgins, *Biometrics: Identity Assurance in the Information Age*, New York, NY: McGraw-Hill, 2003.

[13] Y. Du, B. Bonney, R. W. Ives, D. M. Etter, and R. Schultz, “Partial Iris Recognition Using a 1-D Approach: Statistics and Analysis,” presented at the *2005 IEEE International Conference on Acoustics, Speech, and Signal Processing* (ICASSP), Philadelphia, PA, Mar 2005.

[14] Y. Du, R.W. Ives, D.M. Etter, T.B. Welch, and C.-I Chang, "One dimensional approach to iris recognition", *Proceedings of the SPIE*, pp. 237-247, Apr., 2004.