

Effect of Severe Image Compression on Iris Recognition Performance

John Daugman and Cathryn Downing

Abstract—We investigate three schemes for severe compression of iris images in order to assess what their impact would be on recognition performance of the algorithms deployed today for identifying people by this biometric feature. Currently, standard iris images are 600 times larger than the IrisCode templates computed from them for database storage and search; but it is administratively desired that iris data should be stored, transmitted, and embedded in media in the form of images rather than as templates computed with proprietary algorithms. To reconcile that goal with its implications for bandwidth and storage, we present schemes that combine region-of-interest isolation with JPEG and JPEG2000 compression at severe levels, and we test them using a publicly available database of iris images. We show that it is possible to compress iris images to as little as 2000 bytes with minimal impact on recognition performance. Only some 2% to 3% of the bits in the IrisCode templates are changed by such severe image compression, and we calculate the entropy per code bit introduced by each compression scheme. Error tradeoff curve metrics document very good recognition performance despite this reduction in data size by a net factor of 150, approaching a convergence of image data size and template size.

Index Terms—Biometrics, image compression, image segmentation, iris recognition, JPEG2000, region of interest, receiver operating characteristic (ROC) curves.

I. INTRODUCTION

DATA compression is one of several disciplines rooted in information theory having relevance to biometric technologies for identifying persons, and its significance extends beyond the practical matter of data-storage requirements. One of Shannon's fundamental insights in formulating information theory [1] was that the entropy of a random variable measures simultaneously its information content (expressed in bits) and its compressibility without loss (to the same number of bits). This link between entropy, informativeness, and compressibility extends also to other measures that apply to biometrics. For example, the relative entropy between two distributions is one way to measure how well a biometric technique separates samples from the same versus different people. The amount of variability in a given biometric across a population, or in different samples from the same source, is also captured by conditional entropies, with a larger entropy signifying greater randomness. Finally, the similarity between pairs of biometric

templates is reflected by their mutual information: the extent to which knowledge of one sample predicts the other. All of these properties are deeply connected with the compressibility of biometric data.

An extreme variant of Shannon's insight was expressed by Kolmogorov [2] in his notion of minimal description length, which defined the complexity of a string of data as the length of the shortest binary program that could generate the data. Creating that program "compresses" the data; executing that program "decompresses" (generates) the data. Fractal image compression is based on this idea; and a data string is said to be Kolmogorov incompressible if the shortest program that can generate it is essentially a data statement containing it, so the data is then its own shortest possible description. Within biometrics, this notion has appeared implicitly under a different rubric in work on synthetic biometrics, seeking methods for artificially synthesizing a biometric image that is indistinguishable in practice from some actual biometric image. Pioneering work in this direction was done by Terzopoulos and Waters [3] for facial images and sequences, by Cappelli *et al.* [4] for fingerprints, and by Cui *et al.* [5] and by Zuo *et al.* [6] for iris images. In the future, such programs for generating particular biometric images might therefore serve as ways to "compress" them in Kolmogorov's sense; and one might even anticipate biometric recognition by comparison of the synthesizing programs. In this work, we investigate a question somewhat related to Kolmogorov's query about the relation between data length and its descriptive length. We ask how severely can raw image data be compressed without significantly affecting the biometric templates computed from the data, or thereby the computation of identity. In the case of iris recognition, we demonstrate compressibility of image data down to the order of size of the standard template itself, with minimal impact on performance.

For reasons more mundane and related to policy conformance, data compression in biometrics is also important because governments, regulatory bodies, and international standards organizations often specify that biometric data must be recorded and stored in image form, rather than in (or in addition to) postprocessed templates that may depend on proprietary algorithms. The reasons are to try to achieve interoperability and vendor neutrality. Enrolling or storing raw image data also makes such deployments and databases more "future proof" since they can benefit from inevitable future improvements in recognition algorithms, simply by enrolling anew the raw data. Finally, a directive for standards bodies, such as the International Organization for Standardization (ISO) [7] and industry consortia, such as the Registered Traveller Interoperability Consortium (RTIC) [8] that embed biometric data into smart cards, is to avoid incorporating patented techniques into

Manuscript received July 24, 2007. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Vijaya Kumar Bhagavatula.

The authors are with The Computer Laboratory, University of Cambridge, Cambridge CB3 0FD, U.K. (e-mail: john.daugman@CL.cam.ac.uk; cathryn.downing@CL.cam.ac.uk).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TIFS.2007.916009

data formats and standards, as that would effectively confer monopolies. But storing images instead of templates can imply almost a thousand-fold increase in data size, with consequences such as greatly increased data transmission times and inability to embed the image data in the allocated space in smart cards, which in the case of the RTIC specification [8] for iris images is a mere 4000 bytes per eye. Hence, questions of compressibility and about the effects of lossy image compression on recognition performance become critical.

In summary, both for fundamental scientific reasons related to information theory, and also for practical reasons related to standards, data formats, and storage media, it is important to ask: How much raw image data is really needed for biometric recognition technologies to perform effectively? A watershed event in fingerprint technology occurred in 1993 when the FBI adopted the wavelet scalar quantization (WSQ) protocol [9] to compress vast libraries of fingerprint photograph cards that were digitized to 500 dpi, previously stored in acres of filing cabinets, to achieve compression ratios of typically 10:1 or 15:1. In the relatively new field of iris recognition [10], [11], a pioneering study of iris compressibility was undertaken by Rakshit and Monro [12] showing that if segmented and normalized iris data were extracted in polar form, this “unwrapped” polar data structure could be compressed to 2560 bytes or even less without impairing recognition performance. In this paper, we investigate three compression schemes that retain the native rectilinear image array format but compress it to as little as 2000 bytes while still allowing very good recognition performance on the difficult Iris Challenge Evaluation (ICE-1) iris database available from the National Institute of Standards and Technology (NIST) [13]. We also document interoperability between those images subjected to the compression regimes and their uncompressed form, and we find that on average only 2% to 3% of the bits within the computed 512-byte iris templates (“IrisCodes”) are affected even when the net image reduction factor reaches 150:1.

II. SIMPLE CROPPING AND JPEG COMPRESSION

An obvious first step to reduce the image data size from the standard iris image format of 640×480 pixels with 8-bit grayscale data per pixel, consuming 307 200 bytes, is to crop the image to a smaller region containing the iris, and then to JPEG compress this cropped image. We ran the eye-finding part of the standard algorithms [11] that are used in all current public deployments of iris recognition, on all images in the publicly available NIST [13] ICE1Exp1 database, which contains 1425 iris images from 124 subjects with “ground-truth” information given about which images were taken from the same iris. This database contains many images in which the iris is partly outside the full (640×480) image frame, or is severely defocused, occluded by eyelids or printed patterned contact lenses, interlace corrupted, or with the gaze of the eye directed away from the camera. The real-time algorithms for iris finding and encoding at video rates (30 frames/s) have been described before in detail [11] and will not be reviewed again here. The algorithms correctly localized the iris in all images and produced from each one a new cropped image of 320×320 pixels with the iris centered in it. For those images

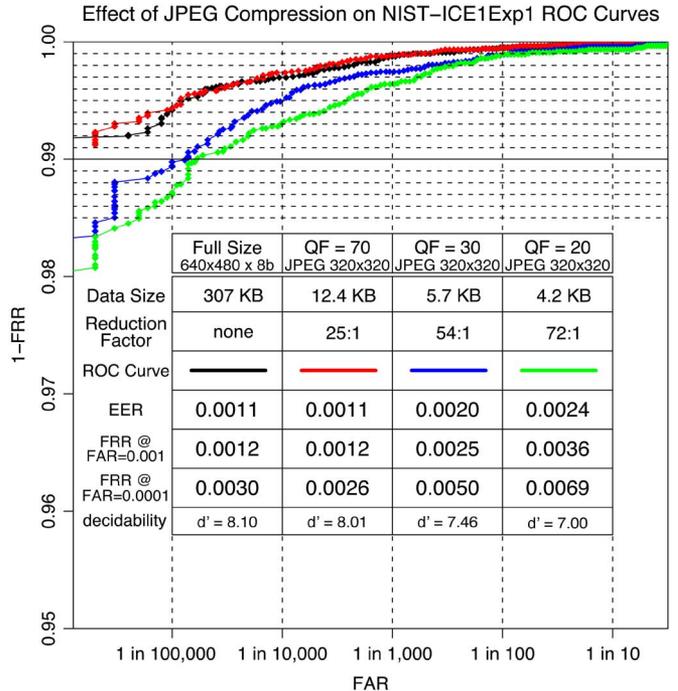


Fig. 1. ROC curves in semilogarithmic coordinates for the NIST [13] ICE1Exp1 iris database, showing the impact of simple data reduction methods on performance. The black curve shows baseline performance on the original database of full-size images. The red curve shows the effect of simple cropping to 320×320 pixels after automatically locating and centering each iris, followed by JPEG compression at $QF = 70$. The blue and green curves show the effects of more severe JPEG compression at $QF = 30$ and $QF = 20$.

in which the iris was partly outside the original image frame, the missing pixels were replaced with black ones. For those in which the algorithms detected that the gaze was directed away from the camera, as gauged by projective deformation of the eye shape, a corrective affine transformation was automatically applied, which effectively “rotated” the eye in its socket into orthographic perspective on-axis with the camera. The new gallery of 1425 cropped and centered iris images was then JPEG compressed [14], [15] by various factors using the Linux tool `cjpeg` with several specified quality factors (QF). These lossily compressed galleries were then decompressed using Linux tool `djpeg` to recover image arrays. The standard algorithms were then run on all of the decompressed images to relocalize the iris, generate IrisCodes (phase bit sequences) [11], and then compare each IrisCode with all others, from both the same and different eyes, to measure the loss in recognition performance against baseline performance for the same original (uncompressed and uncropped) images.

Biometric recognition performance is usually measured by generating receiver operating characteristic (ROC) curves, which plot the tradeoff between two error rates [false accept rate (FAR) and false reject rate (FRR)] as the decision threshold for similarity scores is varied from conservative to liberal. It is common to tabulate specific points on such tradeoff curves, such as the FRR when the decision threshold causes an FAR of 1 in 1000 or of 1 in 10 000, and the point at which the two error rates are equal, $FRR = FAR = EER$, the equal error rate. Such ROC curves and tabulations are presented in Fig. 1

for the NIST [13] ICE-1 gallery, both for baseline performance (uncompressed and uncropped: black curve) and for three JPEG quality factors (colored curves). The coordinates for the ROC curves are semilogarithmic: the ordinate plots 1-FRR linearly, over just the upper 5% of its possible range, while the abscissa logarithmically spans many factors of 10 in FAR, to nearly as low as one in a million. The number of images and the mix of subjects in this NIST iris database allows 12 214 same eye matches to be tested, and it allows 1 002 386 different eye comparisons to be done, which means that one cannot measure an FAR between 0 and 1 in a million; this determines the limit of the ROC curves on the left extreme of these graphs.

The JPEG quality factors (QF) studied here were 70, 30, and 20, producing cropped image file sizes averaging 12 400 bytes, 5700 bytes, and 4200 bytes (red, blue, and green ROC curves, respectively). Including the initial three-fold reduction in file size due merely to cropping the images to 320×320 pixels, these net data reduction factors relative to the original full-size images therefore average 25:1, 54:1, and 72:1, respectively. The red ROC curve in Fig. 1 shows that at a JPEG quality factor of 70 and an overall data reduction factor of 25:1, no performance loss relative to the baseline (black) ROC curve is detectable. (Indeed there is even some suggestion of a small benefit from compression, possibly due to denoising.) The blue and green ROC curves show that for this scheme based only on image cropping and JPEG compression, using a QF in the range of 20 to 30 produces image file sizes in the range of 5000 bytes but at the cost of roughly doubling the FRRs and EER.

Clearly, one could do better by a form of cropping which extracted only the iris pixels, so that the JPEG compression did not waste bytes on noniris pixels. Iris templates are usually computed from a polar or pseudopolar coordinate mapping of the iris, after locating its inner and outer boundaries. The ISO/IEC 19794-6 Iris Image Data Standard [7] specifies two optional methods of sampling iris pixels in polar coordinates so that image data (pixels) rather than computed templates (which would generally be proprietary) could be used for an interoperable data interchange. However, both methods suffer from the fact that polar mappings depend strongly upon the choice of origin of coordinates, which may be prone to error, uncertainty, or inconsistency. Unlike rectilinear coordinates, for which a shift error has no more effect than a shift, in polar mappings a shift error in the choice of coordinate origin can cause large distortions in the mapped data, with no way to recover from such deformations.

In one of the optional polar methods (6.3.2.3) of the standard [7], the mapping extends from the determined center of the pupil to some distance beyond the outer boundary of the iris. Unfortunately, whatever fraction of this diameter is the pupil diameter (typically about 40%), that same fraction of the data is wasted on encoding the black pixels of the pupil, since it is a polar grid. In the other optional polar method (6.3.2.2), circular models are assumed for the inner and outer boundaries of the iris, and the image data is mapped just between those. But in fact for many irises these boundaries cannot be well described as circles; two examples are shown in Fig. 2. In the lower left corner of each picture are two wavy “snakes;” the lower snake is the curvature map of the pupil boundary, and the upper snake is the curvature

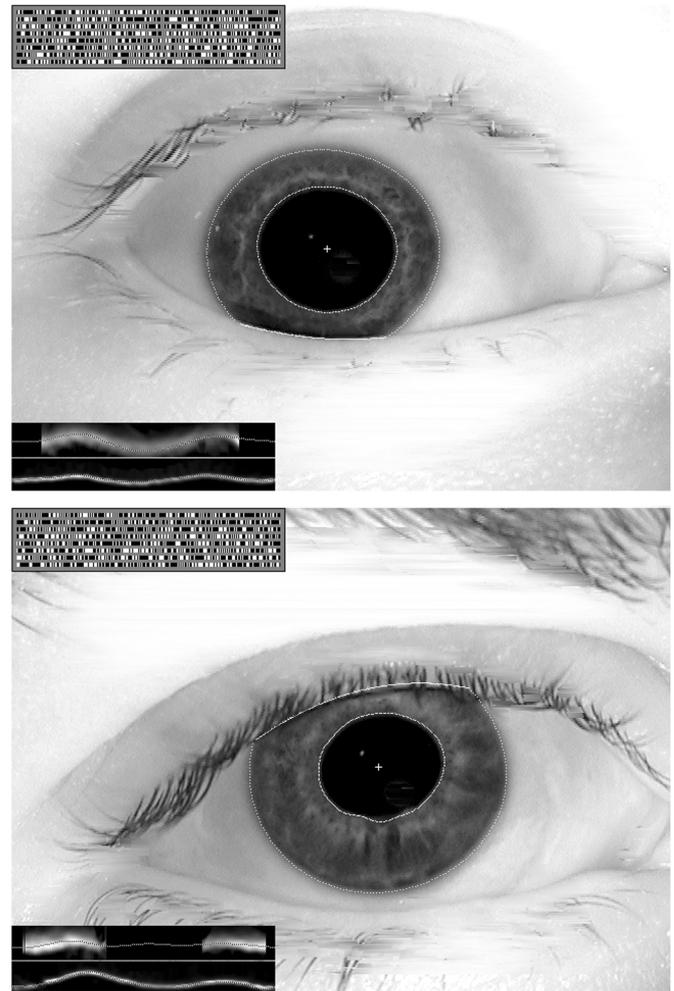


Fig. 2. Many irises have noncircular boundaries, creating problems for polar mappings. The box in the lower left of each image shows the inner and outer boundary curvature maps, which would be flat and straight if they were circles. Active contours enhance iris segmentation and enable flexible coordinate systems; the dotted curves are Fourier series approximations. The bit streams shown in the upper left are the computed IrisCodes.

map of the iris outer boundary. If the assumptions of circular boundaries were valid these should both be straight lines, corresponding to a constant radius of curvature. Clearly they are not. Instead, the dotted curves shown fitting the data, both along the actual iris boundaries and as the skeleton of each snake, are Fourier series expansions of the boundaries using up to 16 Fourier components. (The dc term in such Fourier series expansions corresponds to a simple circular model, and this value is its radius.) Such flexible “active contours” are very important for achieving good iris mappings, but they are not consistent with the polar mappings specified in the data format standard [7]. So we seek a compressible data format that retains rectilinear coordinates, thereby avoiding the problems with polar mappings mentioned before, but in which the iris data alone receives nearly all of the coding budget.

III. REGION-OF-INTEREST (ROI) SEGMENTATION

The standard lossy JPEG coding scheme [14], [15] effectively allocates bytes on an “as needed” basis, meaning that the cost of encoding uniform regions of an image is almost nil, whereas

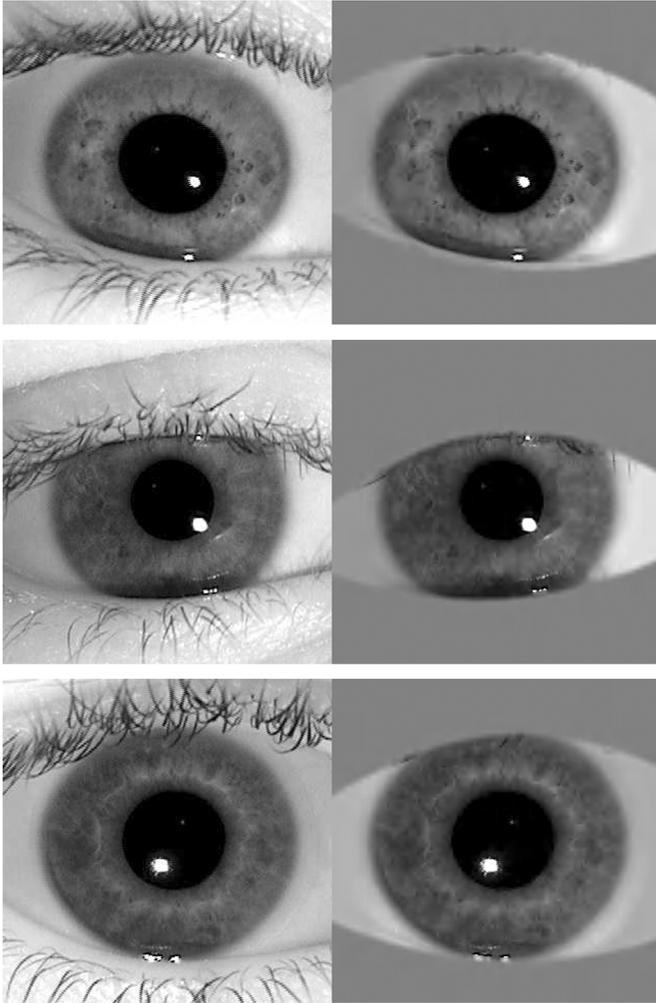


Fig. 3. Region-of-interest isolation of the iris within rectilinear image array formats, to achieve greater compression. Substitution of noniris regions by uniform gray levels prevents wasting the coding budget on costly irrelevant structures such as eyelashes.

image areas containing busy textures, such as eyelashes, may consume much of the available information budget. In uniform regions, the only nonzero discrete cosine transform (DCT) coefficient in each block of 64 frequency components that encode an 8×8 pixel block (data unit) is the dc coefficient specifying their average gray value; all other coefficients are 0 if the data unit is a truly uniform region, or else become 0 after lossy quantization, and so their cost in the zeroes run-length coding stage is essentially nil. Therefore, JPEG encoding of iris images can be made much more efficient if all noniris parts of the image are replaced with a uniform gray value. Such a substitution of pixel values within what is still a rectilinear image array is preferable, from the viewpoint of standards bodies, than actual extraction and mapping of pixel data from a normalized (“unwrapped”) iris because it is desirable to be as shape agnostic and as algorithm neutral as possible. This original rectilinear format is also preferable mathematically because pixels retain constant size and spacing, rather than suffering the polar size distortions and shift sensitivity of unwrapping methods.

JPEG coding schemes lend themselves well to the region-of-interest (ROI) differential assignment of the coding budget. Indeed the JPEG2000 standard [16]–[18], and even the Part 3 extension of the old JPEG standard [14], [15] support variable quantization for explicitly specifying different quality levels for different image regions. In JPEG2000, the MAXSHIFT tool allows specification of an ROI of arbitrary shape. This was explored for biometric face recognition by Hsu and Griffin [19], who demonstrated that recognition performance was degraded by no more than 2% for file sizes compressed to the range of 10 000–20 000 bytes with ROI specification. We now investigate how much compression of iris images can be achieved with minimal impact on iris recognition performance, using the ROI idea without “unwrapping” the iris but retaining a rectilinear array format.

Noniris regions must be encoded in a way that distinguishes sclera from eyelids or eyelashes regions, so that postcompression algorithms can still determine both types of iris boundaries. Therefore we use two different substitution gray levels: a darker one signifying eyelids and a brighter one for the sclera, computed as an average of actual sclera pixels and blending into actual sclera pixels near the iris outer boundary. Since the substitution gray levels are uniform, their coding cost is minimal and could be further reduced by using larger data units. Examples of such ROI segmentation within the rectilinear image array format are shown in the second column of Fig. 3; the first column shows each eye before ROI isolation. The eyelid boundaries were automatically detected by the standard algorithms [11] as the basis for pixel substitution, and the transition to eyelid substitution regions was locally smoothed by a (5×7) kernel to minimize the boundary’s impact on the coding budget. For any given specified QF, the result of iris ROI isolation is typically a two-fold reduction in file size while maintaining a simple rectilinear image format and easy localization of eyelid boundaries in later stages.

The distribution of image file sizes after JPEG compression under various quality factors, with and without ROI segmentation, is shown in the histograms of Fig. 4. In each of the six schemes shown, the range of file sizes obtained spans a factor (max/min) of about 3:1. This unpredictability in the actual file size that will be obtained when specifying a given QF for JPEG compression is disadvantageous in biometric data storage and transmission schemes that allocate a fixed payload space [8]. However, for each QF studied, the benefit of the ROI iris isolation is clear: it reduces file sizes on average by another factor of two.

The impact of the ROI isolation and file size reduction on iris recognition performance is gauged by the ROC curves in Fig. 5. These show that for each QF studied, iris recognition performance remained about the same as before the ROI isolation (Fig. 1), yet with an achievement of a further two-fold reduction in image data size, even down to the range of just 2000–3000 bytes per image.

IV. JPEG2000 COMPRESSION WITH ROI SEGMENTATION

In 2000, a more powerful version of JPEG coding offering more flexible modes of use, and typically achieving 20–30% further compression at any given image quality, was enshrined

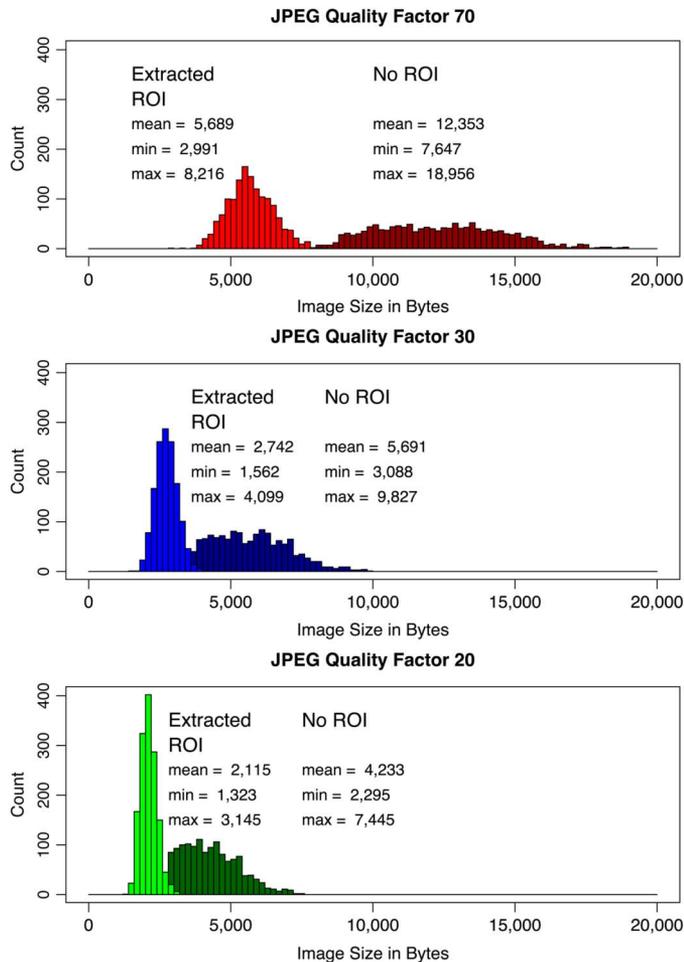


Fig. 4. Distributions of file sizes for the 1425 iris images in the NIST [13] ICE1Exp1 database when JPEG compressed with quality factors of 70, 30, and 20, with and without the ROI isolation of the iris. At every QF there is a clear benefit from the ROI isolation, amounting to typically a factor of two in further file size reduction. Iris recognition performance for each of these six cases is given by the corresponding ROC curves in Figs. 1 and 5.

as the JPEG2000 Standard [17], [18]. Mathematically based on a discrete wavelet transform (DWT) onto Daubechies wavelets rather than the DCT, JPEG2000 does not suffer as badly from the block quantization artifacts that bedevil JPEG at low bit rates, which are due to the fact that the DCT simply chops cosine waves inside box windows with obvious truncation consequences when they are sparse and incomplete. Moreover, the different levels within the multiresolution DWT wavelet decomposition allow local areas within each image tile to be encoded using different subbands of coefficients [18] as needed. The net superiority of JPEG2000 over JPEG in terms of image quality is especially pronounced at very low bit rates, corresponding to severe compression, as we study here, in the range of 0.15 bits per pixel (bpp). Finally, JPEG2000 allows the use of a mask to specify an ROI of arbitrary shape to control the allocation of the encoding budget.

Several mechanisms exist within JPEG2000 for heterogeneous allocation of the coding budget, including tile definition, code-block selection allowing different DWT resolution levels in different tiles, and DWT coefficient scaling. In this paper, we do not explicitly control those parameters nor use the

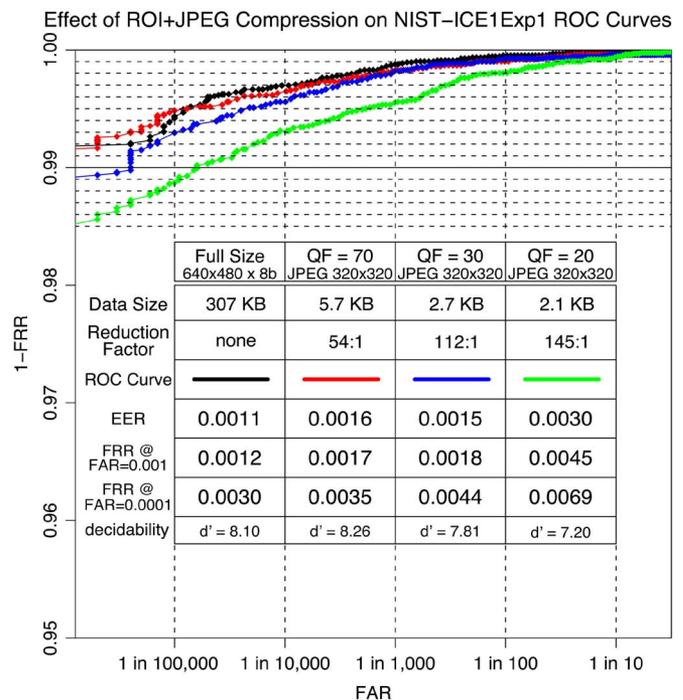


Fig. 5. ROC curves and data-size statistics showing the consequences of ROI isolation before JPEG image compression, so that the available information budget is allocated almost entirely to the iris texture itself. The same quality factors were specified as in the corresponding curves of Fig. 1, and the recognition performance is generally comparable, but now the data reduction factors achieved in each case are twice as great.

MAXSHIFT tool, but rather we use the same pixel substitution method for ROI as described earlier, for comparison purposes. The Linux tools we used for JPEG2000 compression and decompression at various quality factors to document effects on iris recognition performance were `pamtojpeg2k` and `jpeg2ktojam` from the JasPer JPEG2000 and Netpbm libraries. Examples of the resulting ROI+JPEG2000 images can be seen in the second column of the previous Fig. 3 which was used to introduce the ROI method. Those three images were created with a JPEG2000 compression factor (CF) of 50 and thus have a file size of only about 2000 bytes. Whereas JPEG generates widely varying file sizes to deliver any given QF, as was seen in the histograms of Fig. 4, JPEG2000 creates file sizes that are closely predictable from the specified CF. In our experience of compressing several thousand iris images with JPEG2000, the standard deviation of the distribution of resulting file sizes was usually only about 1.6% of the mean, for any given CF. (This variation is narrower even than the width of a bin in the Fig. 4 histograms.) Predictable file size is an important benefit for fixed payload applications [8].

Starting with the same gallery of cropped (320×320) and ROI-isolated iris images illustrated in Fig. 3 that led to the ROC curves of Fig. 5 after JPEG compression at various QF values, we created new galleries compressed by JPEG2000 at CF values of 20, 50, and 60. These galleries had image data sizes of about 5100, 2000, and 1700 bytes, respectively. Fig. 6 presents the ROC curves that the galleries generated, together for comparison with the black ROC curve for the baseline gallery (uncropped, uncompressed, not ROI isolated). It is clear that com-

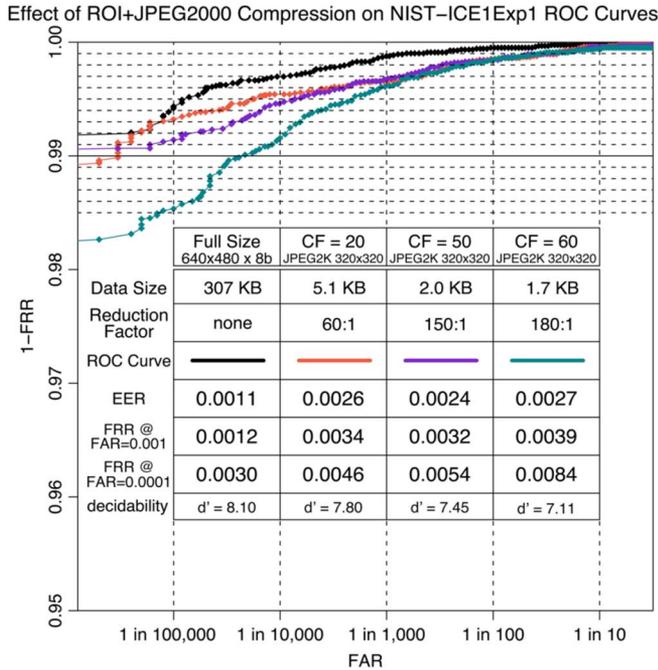


Fig. 6. ROC curves and data-size statistics showing iris recognition performance when the cropped and ROI-isolated images are compressed using JPEG2000 at various compression factors. Performance with file sizes of merely 2000 bytes ($CF = 50$, purple curve) remains remarkably unimpaired compared to a baseline (black curve); but further compression begins to exact a high toll (blue-green curve).

pression as severe as 0.156 bpp ($CF = 50$, file size 2.0 kB, purple curve) still allows remarkably good iris recognition performance. For example, the FRR remains below 1% at an FAR of 1 in 100 000. It seems extraordinary that image arrays reconstructed from as little as 2000 bytes of data are still so serviceable for iris recognition. It is possible that part of the explanation lies in the similarity between the Daubechies wavelets used for the DWT in JPEG2000 coding, and the Gabor wavelets used in our creation [11] of the IrisCode itself, so that information lost in such severe compression is not used in the IrisCode anyway. However, a watershed seems to exist at 2000 bytes, since a pronounced degradation becomes evident when images are further compressed to 1700 bytes ($CF = 60$, blue-green ROC curve in Fig. 6).

V. COMPARING THE EFFECTS OF THE COMPRESSION SCHEMES

In this paper we have focused on ROC curves, which reflect the overlapping tails of the two distributions of similarity scores computed for images from same or different eyes. The similarity score is a normalized Hamming distance (HD), which is the fraction of bits disagreeing between two IrisCodes among the bits compared [11]. It is informative to see the full distributions of HD scores, which we present in Fig. 7 for two of the compression schemes. In each panel, two different ordinate axis scales are used to facilitate visualization since there are 1 002 386 counts in the “all against all other” distribution (magenta) created by comparing different eyes, but only 12 214

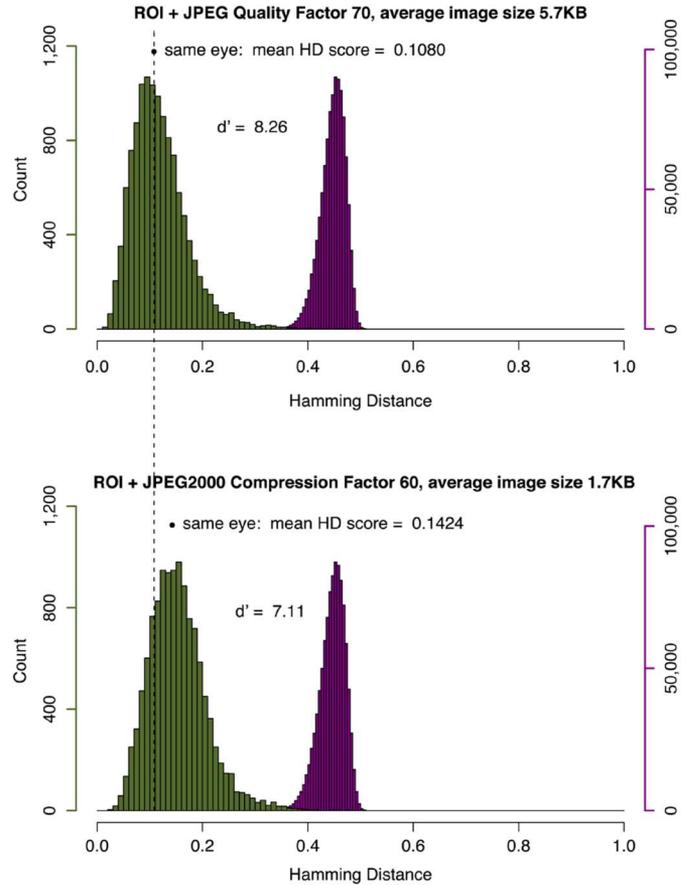


Fig. 7. Distributions of Hamming distance scores comparing the same and different eyes in the NIST database, for two of the image compression schemes bracketing the range of schemes studied. Even in the most severe case (lower panel) using images compressed to only 1700 bytes, the dual distributions have little overlap and so decisions about identity remain robust.

counts in the distribution (olive) made by all same eye comparisons across the database. The upper panel shows the distributions obtained with ROI+JPEG compression at $QF = 70$, which created an average file size of 5700 bytes and generated the red ROC curve in Fig. 4. As was evident in Fig. 4, the recognition performance obtained with that compression scheme was almost indistinguishable from the baseline performance (black ROC curve: no compression, ROI, or cropping). The dual distributions for that baseline case are likewise indistinguishable from the upper panel in Fig. 7, as one would expect, and so we do not include them here. The lower panel shows the distributions obtained with ROI+JPEG2000 compression at $CF = 60$, which created an average file size of just 1700 bytes and generated the blue-green ROC curve in Fig. 6. It is remarkable that such extremes of compression do not have catastrophic effects on the separability of the pair of distributions. Instead, we see in Fig. 7 that the distribution obtained from different eyes (magenta) is virtually unchanged, whereas the distribution obtained from same eye images (olive) is shifted to the right by a small amount, corresponding to an increase in the mean HD score from 0.1080 to 0.1424 as indicated by the two dots and a projected vertical line for comparison.

Information theory provides certain metrics for defining the “distance” between two random variables in terms of their entire probability distributions. When both random variables are distributed over the same set of possible outcomes, such as the HD scores that were tallied in the histograms for same and for different eyes in Fig. 7, then the relative entropy or Kullback–Leibler distance is a natural way to measure the overall distance between the two distributions. As a measure of separation, it is also called the “information for discrimination.” Unfortunately, this measure becomes undefined if there are some values that only one random variable can have while other values are accessible only to the other random variable. Since the distributions of HD scores obtained from comparisons between different eyes in Fig. 7 vanish for scores smaller than about 0.3, and likewise the score distributions for same eyes attenuate to zero over much of the other distribution, the calculated Kullback–Leibler distance between these distributions is infinite and meaningless, unless based on nonvanishing theoretical models for them or by adding arbitrary quantities that then become decisive for this metric.

An alternative family of distance metrics, encompassing the Fisher ratio and Z-scores, defines distance in terms of the difference between the means of the two distributions, normalized by some function of their standard deviations. One such is the d' metric of decidability in signal detection theory, defined as $d' = |\mu_1 - \mu_2| / \sqrt{0.5(\sigma_1^2 + \sigma_2^2)}$, where μ_1 and μ_2 are the means and σ_1 and σ_2 are the standard deviations. A limitation of this metric is that by considering only the first two moments of the distributions, it makes no explicit use of skew, kurtosis, and higher moments that are more sensitive to mass in the tails. Thus, d' might be said to take a “Gaussian view” of the world, whereas the skewed distributions in Fig. 7 are clearly not Gaussian. Nonetheless, we have included within the ROC graphs in Figs. 1, 5, and 6 the d' scores for each underlying pair of distributions obtained with each of the compression schemes studied. They show a small but systematic trend of deterioration with more aggressive levels of image compression. But as is clear from the two bracketing extremes presented in Fig. 7, the separability of the two underlying distributions remains remarkable, despite the massive compression factor reaching 180:1 reduction from the original images.

Another metric often used in decision theory to summarize overall performance by a single scalar statistic is the area under the ROC curve. Clearly, a value of 1.0 represents perfection since it arises only from the complete absence of overlap between the two distributions. The ten different ROC curves plotted in Figs. 1, 5, and 6 appear to have significant amounts of missing area, but this is an illusion due to the logarithmic abscissa and the magnified ordinate ranging just between 0.95 and 1.00. In fact the area under the baseline black ROC curve (present in all three figures) is 0.999985, and the areas under most of the other nine curves are reduced from this value in only the 6th decimal place.

Finally, it is interesting to compare visually some examples of the iris images after compression to a constant data size of 2000 bytes using the three different schemes. Each column of Fig. 8 is from the same NIST iris image; the rows represent the different schemes. The top row is simple JPEG compression of

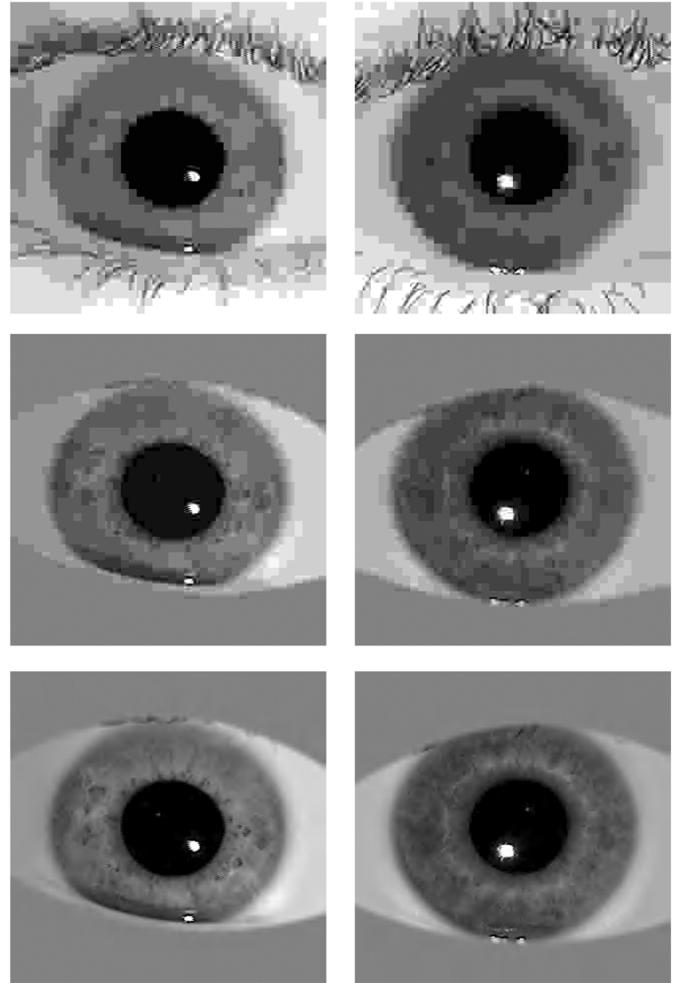


Fig. 8. Visual comparison of the three schemes for iris image compression, for images all compressed to the same data size of 2000 bytes. The left column is NIST 239230; the right is NIST 239343. Top row: simple JPEG compression of the cropped (320×320) images. Middle row: JPEG compression of the cropped images after ROI isolation. Bottom row: JPEG2000 compression of the cropped and ROI-isolated images. Iris recognition performance of this third scheme is shown by the purple ROC curve ($CF = 50$) in Fig. 6.

a cropped (320×320) image but without ROI isolation. Most of the 2000-byte budget is wasted trying to encode eyelashes, and the cost on iris texture is horrendous. The middle row shows improvement after ROI isolation, so most of the JPEG budget is allocated to the iris, but the result is still very poor. The bottom row shows the result of combining the cropping, ROI isolation, and JPEG2000 compression for the same iris images. The improvement is visually remarkable, and it is confirmed by very good iris recognition performance as summarized by the purple ROC curve ($CF = 50$) in Fig. 6.

VI. EFFECT OF SEVERE COMPRESSION ON AN IMAGE-QUALITY METRIC PREDICTING PERFORMANCE

A growing focus of biometric research concerns quality metrics that may predict recognition performance, and that therefore might be used to control the image-acquisition process or to qualify images. It is relevant to examine how some of the compression regimes explored in this paper affect such iris image-quality metrics.

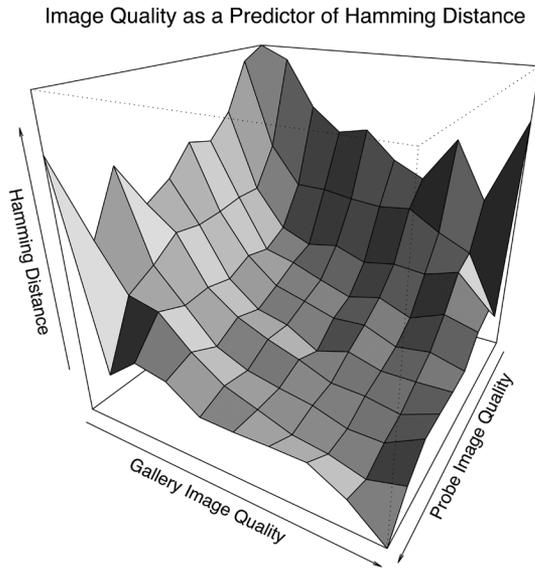


Fig. 9. Effect of image quality on Hamming Distance between same-eye NIST images. The main elements of the quality metric are image focus [11], motion, proportion of iris visible, and number of code bits deemed reliable.

Image Quality as a Predictor of False non-Match Rate

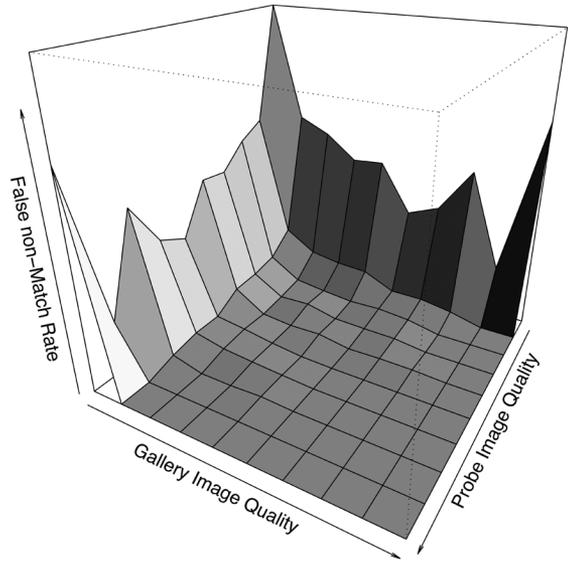


Fig. 10. Image quality as a predictor of false nonmatch rate for same-eye NIST images. The quality metric is here applied to the original uncompressed images, in order to validate its power to predict performance.

For all “probe” and “gallery” pairings of images in the NIST iris database that arise from a given eye and therefore should match, we plotted their match Hamming Distance as a function of both of their quality scores as measured by “Q.” This metric is based mainly on the focus score of each image (computed by 2-D Fourier methods as detailed in [11, Appendix] and on the sharpness of the pupillary boundary as a further indicator of focus), and on detection of interlaced raster shear signifying subject motion during the acquisition of a frame. It is also based on calculating the percentage of the iris that is not occluded by eyelids or eyelashes or by specular reflections, and on the number of bits in its computed IrisCode [11] that are deemed to be significant and reliable. A surface plot showing dependence of match Hamming Distance on both “Q” scores for the available same-eye image pairings in the NIST database is presented in Fig. 9. Clearly, it slopes down almost monotonically as a function of both “Q” scores.

Since the Hamming Distance is the metric on which match decisions are based, the surface plot in Fig. 9 implies directly how the false nonmatch rate (FnMR) is predicted by the image-quality metric. An actual FnMR at a threshold of $HD = 0.32$ is shown in Fig. 10 as another surface plot over joint “Q” scores for these image pairings, and the same data is shown as a contour plot in Fig. 11. Clearly, the quality metric “Q” is a very strong predictor of matching performance.

Therefore, we examine how the three different JPEG2000 compression regimes whose ROC curves were presented in Fig. 6 affect the quality metric “Q” for these images. This is shown in Fig. 12 (using the same color coding as in Fig. 6 for each regime) as a group of overlaid scatter plots. The “Q” scores along the abscissa are for the original uncompressed images after cropping and the region-of-interest isolation, and the scatter plots of “Q” scores plotted as a function of those

Image Quality as a Predictor of False non-Match Rate

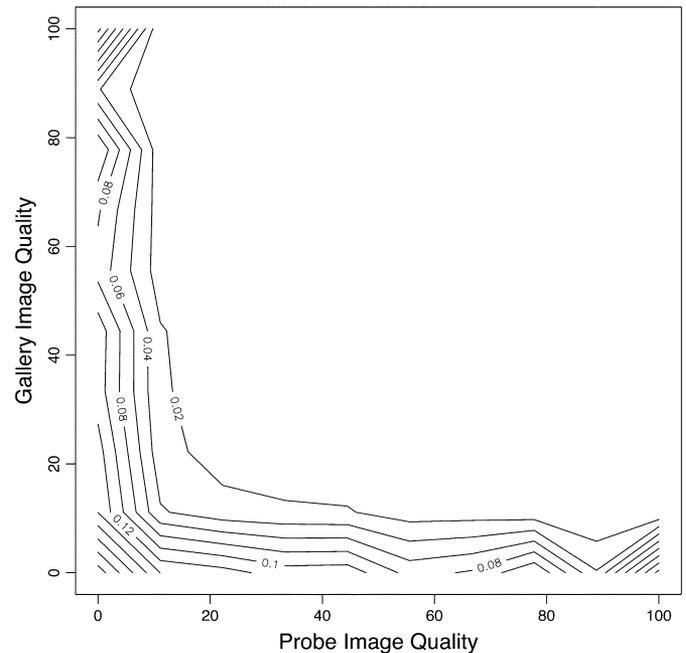


Fig. 11. Effect of image quality on the false nonmatch rate for same-eye NIST images, before application of any of the compression schemes. This baseline contour plot is derived from the surface plotted in Fig. 10.

reveal the impact of each of the three compression regimes ($CF = 20$, $CF = 50$, and $CF = 60$) on this quality metric. The minimal effect of each compression regime on “Q” scores is captured by the correlation coefficients ρ_{CF} between the scores for uncompressed images and the scores after compression by each CF factor. These correlation coefficients are $\rho_{20} = 0.986$, $\rho_{50} = 0.974$, and $\rho_{60} = 0.971$, respectively.

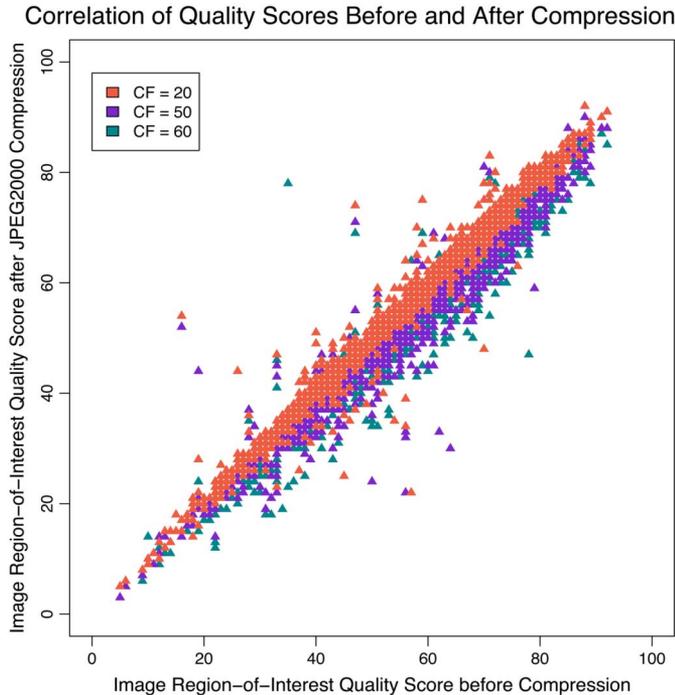


Fig. 12. Scatter plots for quality scores before and after compression under each of the JPEG2000 regimes whose ROC curves were shown in Fig. 6. The color coding for each of the compression factors in these figures is the same.

VII. CONCLUSION

We have studied the effects of three schemes for image compression on iris recognition performance, leading to the surprising conclusion that even images compressed as severely as 150:1 from their original full-size formats, to just 2000 bytes, remain very serviceable. It is important to use region-of-interest isolation of the iris within the image so that the coding budget is allocated almost entirely to the iris; and it is important to use JPEG2000 instead of JPEG as the compression protocol. Advantages of this overall approach from the perspective of standards bodies and interoperability consortia are that the compact image data (when decompressed) is a native rectilinear array; no proprietary methods are required; and the distortions that can arise from alternative coordinate transformation methods such as polar unwrapping or polar sampling are avoided.

As concluding measures, we compared the IrisCodes generated under each scheme to those generated for the corresponding original uncompressed images. The entropy $H = -\sum_i p_i \log_2(p_i)$ or uncertainty per code bit caused by each compression scheme is tabulated in Table I. For reference, the entropy associated with the states of bits in IrisCodes calculated from different images of the same eye, due merely to variation in image capture, is typically 0.506 bit; Table I shows that the corrupting effect of the image compression schemes is much less than this native uncertainty in the bits of IrisCodes for a given eye. The final column of Table I tabulates, as interoperability scores, the average HD (fraction of disagreeing bits) between the IrisCodes obtained before and after image compression for each scheme and for each compression parameter. They indicate that only about 2% to 3% of the IrisCode bits change as a consequence of image compression

TABLE I

SUMMARY OF THE COMPRESSION SCHEMES, RESULTING FILE SIZES, AND THEIR EFFECTS ON COMPUTED IRISCODES, EXPRESSED AS ENTROPY PER CODE BIT AND AS THE FRACTION (HD) OF BITS THAT WERE CHANGED FROM THOSE COMPUTED FOR THE ORIGINAL FULL-SIZE IMAGES

Strategy	Compression Parameter	Average Image Size	Entropy / code bit	Interoperability HD
Cropping (320×320) + JPEG Compression	QF = 70	12.4 kB	0.053 bit	0.006
	QF = 30	5.7 kB	0.087 bit	0.011
	QF = 20	4.2 kB	0.147 bit	0.021
Cropping + ROI + JPEG Compression	QF = 70	5.7 kB	0.112 bit	0.015
	QF = 30	2.7 kB	0.147 bit	0.021
	QF = 20	2.1 kB	0.199 bit	0.031
Cropping + ROI + JPEG2000 Compression	CF = 20	5.1 kB	0.130 bit	0.018
	CF = 50	2.0 kB	0.179 bit	0.027
	CF = 60	1.7 kB	0.219 bit	0.035

even as severe as to 2000 bytes. When considered in the context of Fig. 7 showing the HD distributions for same and different eyes, it is clear that an increment of 0.02 to 0.03 in HD score is a negligible impact indeed. In conclusion, it appears that rough convergence between data length and standard descriptive length for this biometric is possible.

REFERENCES

- [1] C. E. Shannon, "A mathematical theory of communication," *Bell Syst. Tech. J.*, vol. 27, pp. 379–423, 1948.
- [2] A. N. Kolmogorov, "Three approaches to the quantitative definition of information," *Probl. Inform. Transm.*, vol. 1, pp. 4–7, 1965.
- [3] D. Terzopoulos and K. Waters, "Analysis and synthesis of facial image sequences using physical and anatomical models," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 15, no. 6, pp. 569–579, Jun. 1993.
- [4] R. Cappelli, D. Maio, and D. Maltoni, "Synthetic fingerprint-image generation," in *Proc. Int. Conf. Pat. Recog. 15*, 2000, vol. 3, pp. 475–478.
- [5] J. Cui, Y. Wang, J. Huang, T. Tan, and Z. Sun, "An iris image synthesis method based on PCA and super-resolution," in *Proc. 17th Int. Conf. Pattern Recognition*, 2004, vol. 4, pp. 471–474.
- [6] J. Zuo, N. A. Schmid, and X. Chen, "On generation and analysis of synthetic iris images," *IEEE Trans. Inf. Forensics Security*, vol. 2, no. 1, pp. 77–90, Mar. 2007.
- [7] *Information Technology—Biometric Data Interchange Formats—Part 6: Iris Image Data*, Int. Std. ISO/IEC 19794-6, 2005, ISBN 0-580-46456-3.
- [8] Registered Traveler Interoperability Consortium (RTIC), "Technical interoperability specification for the Registered Traveler Program (Version 1.0–Final)," 2006 [Online]. Available: <http://www.rtconsortium.org/>
- [9] J. N. Bradley, C. M. Brislawn, and T. Hopper, "The FBI wavelet/scalar quantization standard for grayscale fingerprint image compression," *Proc. SPIE Applications of Digital Image Processing XIX*, vol. 2847, pp. 293–304, 1996.
- [10] J. Daugman and C. Downing, "Epigenetic randomness, complexity, and singularity of human iris patterns," *Proc. Roy. Soc., B*, vol. 268, pp. 1737–1740, 2001, Biol. Sci..
- [11] J. Daugman, "How iris recognition works," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 14, no. 1, pp. 21–30, Jan. 2004.
- [12] S. Rakshit and D. M. Monro, "An evaluation of image sampling and compression for human iris recognition," *IEEE Trans. Inf. Forensics Security*, vol. 2, no. 3, pp. 605–612, Sep. 2007.
- [13] *National Institute of Standards and Technology, Iris Challenge Evaluation*, [Online]. Available: <http://iris.nist.gov/ice/>.
- [14] G. K. Wallace, "The JPEG still picture compression standard," *Commun. ACM*, vol. 34, no. 4, pp. 30–44, 1991.
- [15] *Information Technology—Digital Compression and Coding of Continuous-Tone Still Images*, Int. Std. ISO/IEC 10918, 1994. [Online]. Available: <http://www.jpeg.org/>.

- [16] A. P. Bradley and F. W. M. Stentiford, "JPEG2000 and region of interest coding," in *Digital Imaging Computing Techniques and Applications*, Melbourne, Australia, 2002. [Online]. Available: <http://www.ee.ucl.ac.uk/~fstentif/DICTA02.pdf>.
- [17] *Information Technology—JPEG 2000 Image Coding System*, Int. Std. ISO/IEC 15444-1, 2004. [Online]. Available: <http://www.itu.int/rec/T-REC-T/en>.
- [18] C. Christopoulos, A. Skodras, and T. Ebrahimi, "The JPEG2000 still image coding system: An overview," *IEEE Trans. Consum. Electron.*, vol. 46, no. 4, pp. 1103–1127, Nov. 2000.
- [19] R. L. V. Hsu and P. Griffin, "JPEG region of interest compression for face recognition," *IDENTIX Document No. RDNJ-04-0102*, 2005.
- John Daugman**, photograph and biography not available at the time of publication.
- Cathryn Downing**, photograph and biography not available at the time of publication.