

Number 685



**UNIVERSITY OF  
CAMBRIDGE**

**Computer Laboratory**

## Effect of severe image compression on iris recognition performance

John Daugman, Cathryn Downing

May 2007

15 JJ Thomson Avenue  
Cambridge CB3 0FD  
United Kingdom  
phone +44 1223 763500  
<http://www.cl.cam.ac.uk/>

© 2007 John Daugman, Cathryn Downing

Technical reports published by the University of Cambridge  
Computer Laboratory are freely available via the Internet:

*<http://www.cl.cam.ac.uk/techreports/>*

ISSN 1476-2986

# 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 persons 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 government database of iris images. We show that it is possible to compress iris images to as little as 2 KB 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. Standard performance metrics such as error trade-off curves 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.

## 1 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 same versus different persons. 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 larger entropy signifying greater randomness. Finally, the similarity between pairs of biometric templates may be measured as their mutual information, also called equivocation entropy: the extent to which knowledge of one sample determines or 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 the present work, we investigate questions related to Kolmogorov’s concept by asking how severely the raw image data can be compressed without significantly affecting the biometric templates computed from the data; and we find a rough convergence of template and image data sizes with minimal impact on iris recognition 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 a raw form, rather than in (or in addition to) post-processed 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 like ISO [7] and industry consortia such as RTIC [8] that embed biometric data into smart cards is to avoid incorporating patented techniques into data formats and Standards, as that would effectively confer monopolies. But storing raw 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 raw data in the allocated space in smart cards, which in the case of the RTIC specification [8] for iris images is a mere 4,000 bytes per eye. Hence questions of compressibility, and about the effects of 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 digitised 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 unimpaired recognition performance for iris data extracted in polar format into data structures of 20,000 bytes (or 0.5 bpp). In this report we document three compression schemes that retain rectilinear image formats but achieve severe compression to as little as 2,000 bytes while still allowing very good recognition performance on the difficult NIST [13] ICE-1 publicly available iris image database. We also document interoperability between all the compression schemes and the uncompressed format, 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.

## 2 Simple cropping and JPEG compression

An obvious first step to reduce image data size from the standard iris image format of  $640 \times 480$  pixels with 8 bits 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 1,425 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 of the full ( $640 \times 480$ ) image frame, or is severely defocused, occluded by eyelids, corrupted by raster shear, aliasing, noise, and motion, 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/sec) have been described before in detail [11] and will not be reviewed again here. The algorithms correctly localised the iris in all images and produced from each one a new cropped image of dimensions  $320 \times 320$  pixels with the iris centered in it. For those images in which the iris was partly outside of 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 back into orthographic perspective on-axis with the camera. The new gallery of 1,425 cropped ( $320 \times 320$ ) 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 lossy image arrays. The standard algorithms were then run on all the decompressed images to re-localise the iris, generate IrisCodes (phase bit sequences) [11], and then compare each IrisCode with all others, from both 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 ROC (Receiver Operating Characteristic) curves, which plot the trade-off between two error rates (False Accept and False Reject Rates, FAR and FRR, also called False Match and False non-Match Rates) as the decision threshold for similarity scores is varied from conservative to liberal. It is common to tabulate specific points on such trade-off curves, such as the FRR when the decision threshold causes an FAR of 1 in 1,000 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 (coloured curves). The coordinates for the ROC curves are semi-logarithmic: 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 1 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 a False Match Rate (or FAR) between 0 and 1 in a million; this determines the limit of the ROC curves on the left extreme of these graphs.

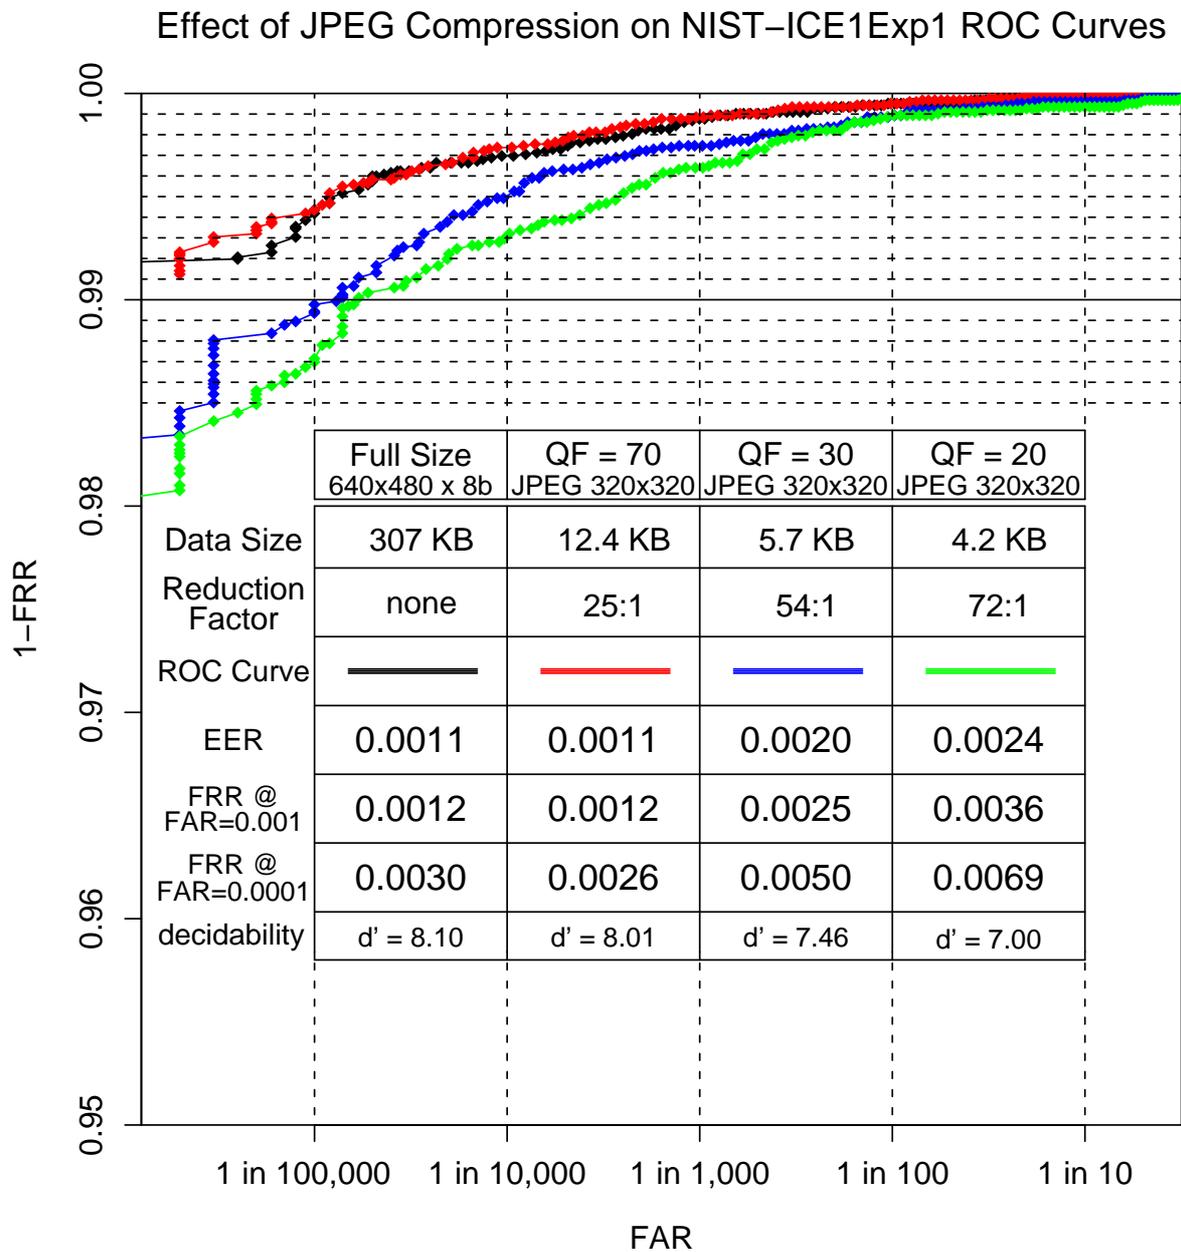


Figure 1: ROC curves in semi-logarithmic coordinates for the NIST [13] ICE1Exp1 iris database, showing the impact of simple data reduction methods on performance. Black curve shows baseline performance on the original database of full-size images. Red curve shows the effect of simple cropping to  $320 \times 320$  pixels after automatically locating and centering each iris, followed by JPEG compression at  $QF = 70$ . Blue and green curves show the effects of more severe JPEG compression at  $QF = 30$  and  $QF = 20$ .

The JPEG quality factors (QF) used here were 70, 30, and 20, producing cropped image file sizes averaging 12,400 bytes, 5,700 bytes, and 4,200 bytes (red, blue, and green ROC curves respectively). Including the initial 3-fold reduction in file size due merely to cropping the images to  $320 \times 320$  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 superimposes almost everywhere onto the black (baseline) ROC curve, indicating that no performance loss is detectable at a JPEG quality factor of 70 and an overall data reduction factor of 25:1. 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 5,000 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 non-iris pixels. Iris templates are usually computed from a polar or pseudo-polar 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 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 both 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 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 also 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 above, but in which the iris data alone receives nearly all of the coding budget.

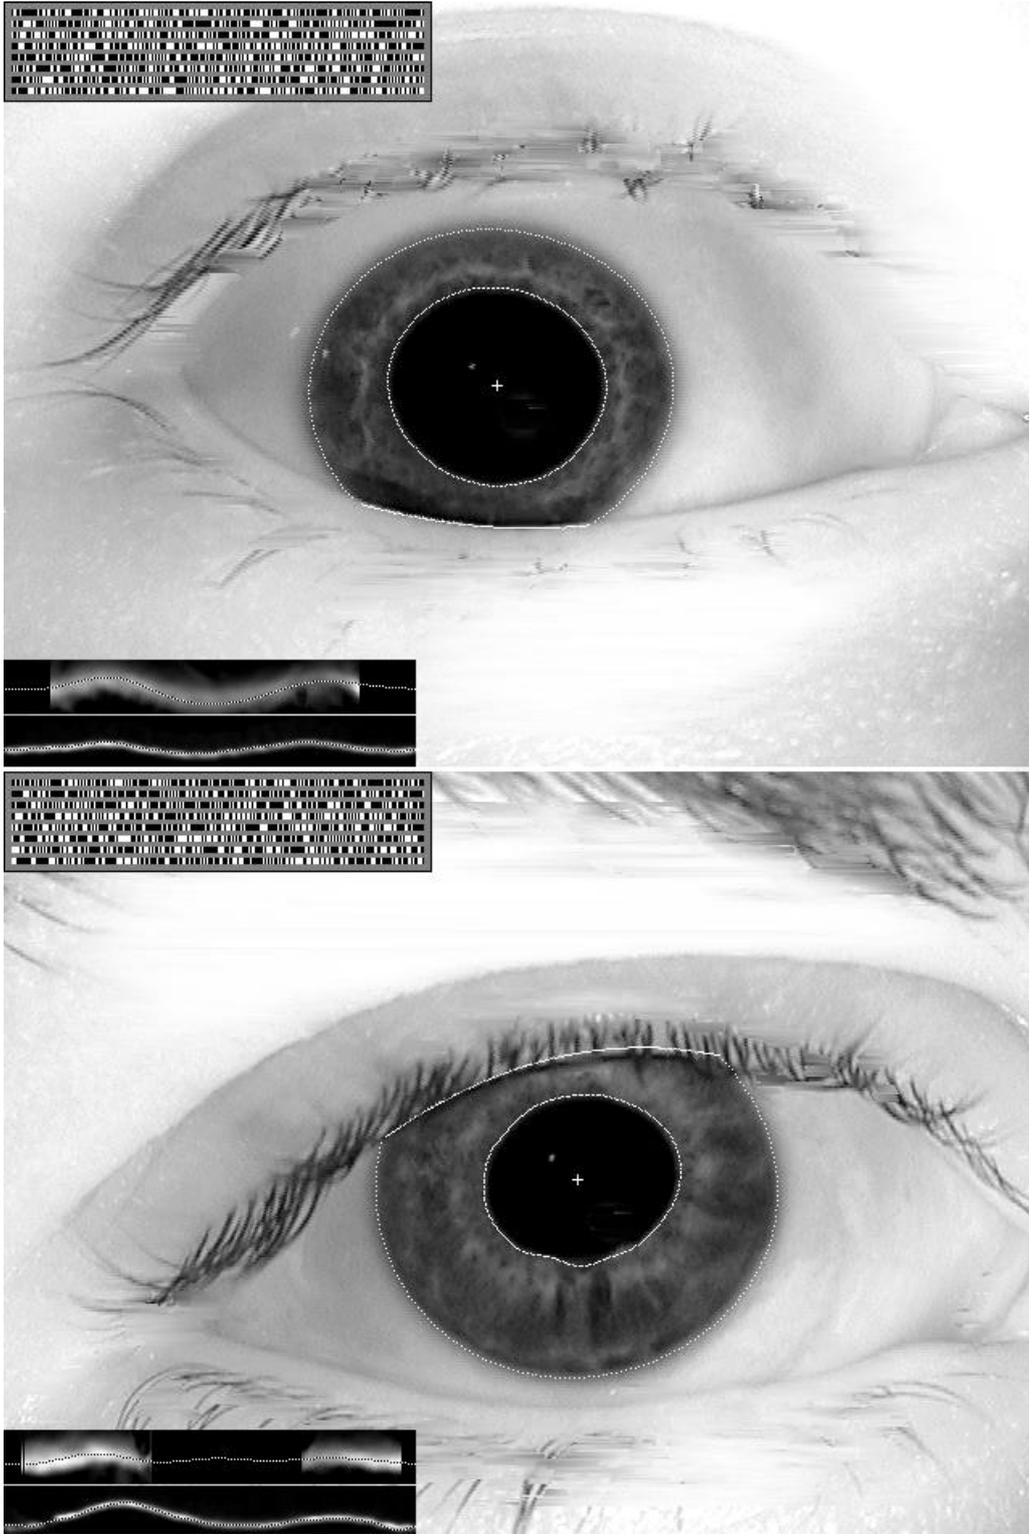


Figure 2: Many irises have non-circular boundaries, creating problems for polar mappings. The box in the lower-left of each image shows the inner and outer boundary curvature maps, which would be flat and straight if they were circles. Active contours enhance iris segmentation and enable flexible coordinate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

### 3 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 image areas containing busy textures such as eyelashes may consume much of the available information budget. In uniform regions, the only non-zero DCT (discrete cosine transform) coefficient in each block of 64 frequency components that encode an 8 x 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 Huffman (run-length) coding stage is essentially nil. Therefore JPEG encoding of iris images can be made much more efficient if all non-iris 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 normalised (“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 region-of-interest (ROI) differential assignment of the coding budget [16]. Indeed the JPEG2000 standard [17, 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. This idea 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 pixel array format for the reasons cited earlier.

Non-iris regions must be encoded in a way that distinguishes sclera from eyelids or eyelashes regions, so that post-compression 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 smoothed by a  $(5 \times 7)$  kernel to minimise 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

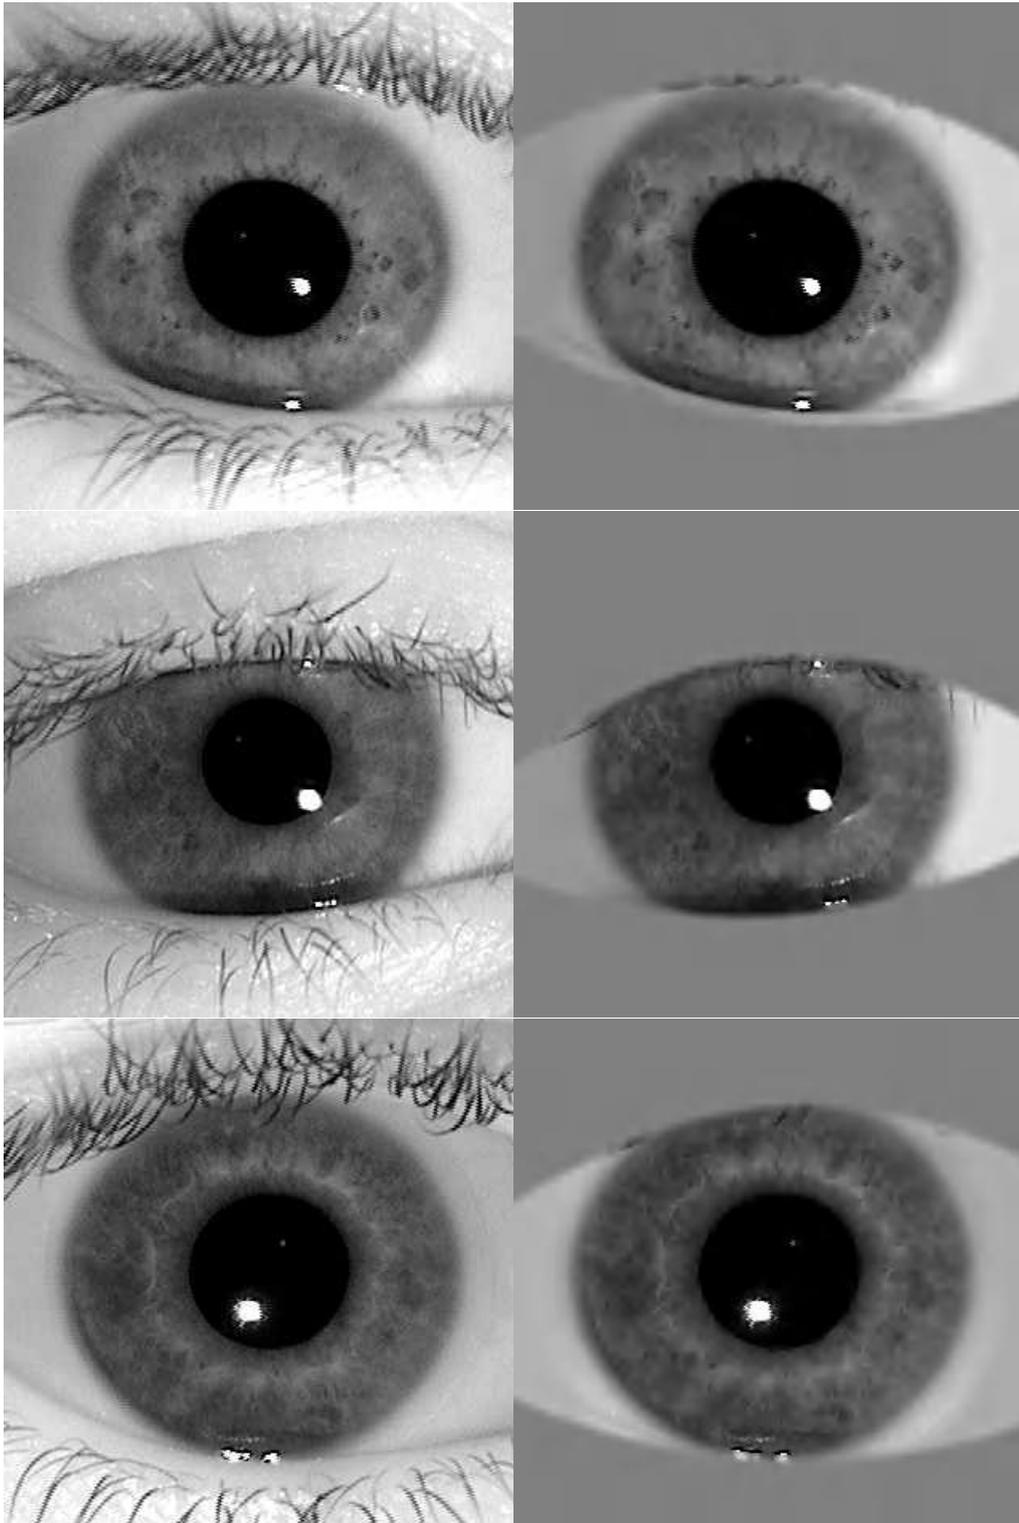


Figure 3: Region-of-interest isolation of the iris within rectilinear image array formats, to achieve greater compression. Substitution of non-iris regions by uniform gray levels prevents wasting the coding budget on costly irrelevant structures such as eyelashes.

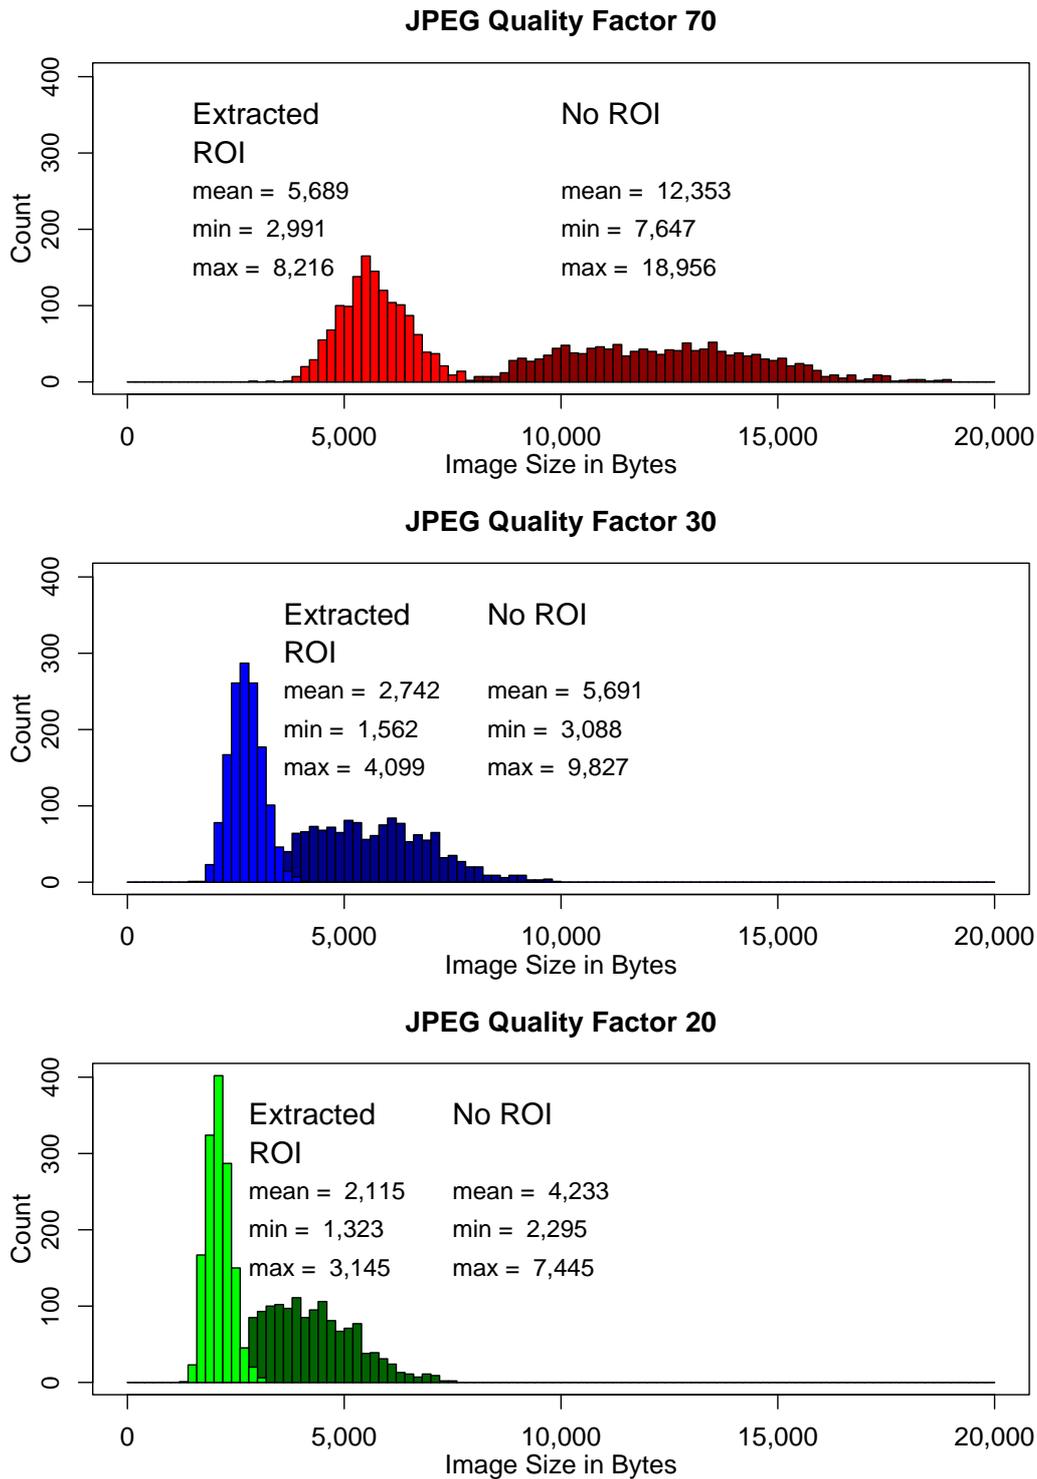


Figure 4: Distribution of file sizes for the 1,425 iris images in the NIST [13] ICE1Exp1 database when JPEG compressed with quality factors of 70, 30, and 20, with and without the ROI isolation of the iris within the rectilinear image array. At every QF there is a clear benefit from the ROI isolation, amounting typically to a factor of two in further size reduction. Iris recognition performance for each of these six cases is given by the corresponding ROC curves in Figs. 1 and 5.

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.

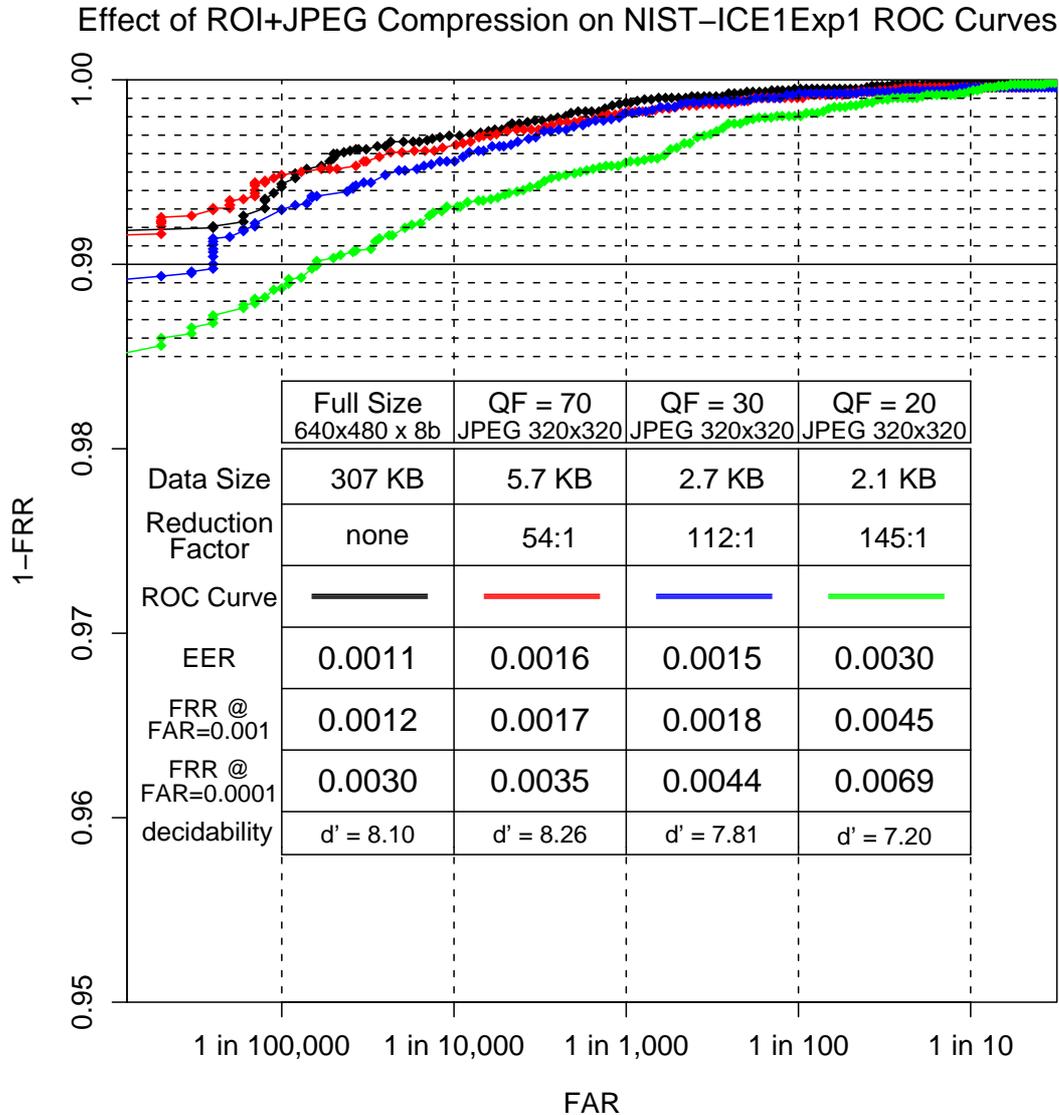


Figure 5: ROC curves and data size statistics showing the consequences of ROI isolation before JPEG image compression, so that the available information budget is allocated almost entirely to the iris texture itself. The same quality factors were specified as in the corresponding curves of Fig. 1, and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

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 achievement of a further two-fold reduction in image data size, even down to the range of just 2,000 – 3,000 bytes per image.

## 4 JPEG2000 compression with ROI segmentation

In 2000 a more powerful version of JPEG coding offering more flexible modes of use, and achieving typically a further 20–30% compression at any given image quality, was enshrined as the JPEG2000 Standard [17, 18]. Mathematically based on a Discrete Wavelet Transform (DWT) onto Daubechies wavelets rather than the Discrete Cosine Transform (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 multi-resolution 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 the very low bit-rates corresponding to severe compression, as we study here. Finally, JPEG2000 allows 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 the present work we do not explicitly control those parameters nor specify an ROI mask, but rather we use the same pixel substitution method for ROI as described above, 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 `jpeg2ktopam` from the `JasPer` JPEG2000 and `Netpbm` libraries. Examples of the resulting ROI+JPEG2000 images can be seen in the second column of the earlier 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 2,000 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 thousands of 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 ( $320 \times 320$ ) 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 5,100, 2,000, and 1,700 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 compression as severe as  $CF = 50$  to a file size of only 2,000 bytes (purple curve) still preserves remarkably good iris recognition performance. For example, the FRR remains below 1% at an FAR of 1 in 100,000. We find it extraordinary that image arrays recovered from as little as 2,000 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 isn't used in the IrisCode anyway. However, a watershed seems to exist at 2,000 bytes, since a pronounced degradation becomes evident when images are further compressed to 1,700 bytes (CF = 60, blue-green ROC curve in Fig. 6).

Effect of ROI+JPEG2000 Compression on NIST-ICE1Exp1 ROC Curves

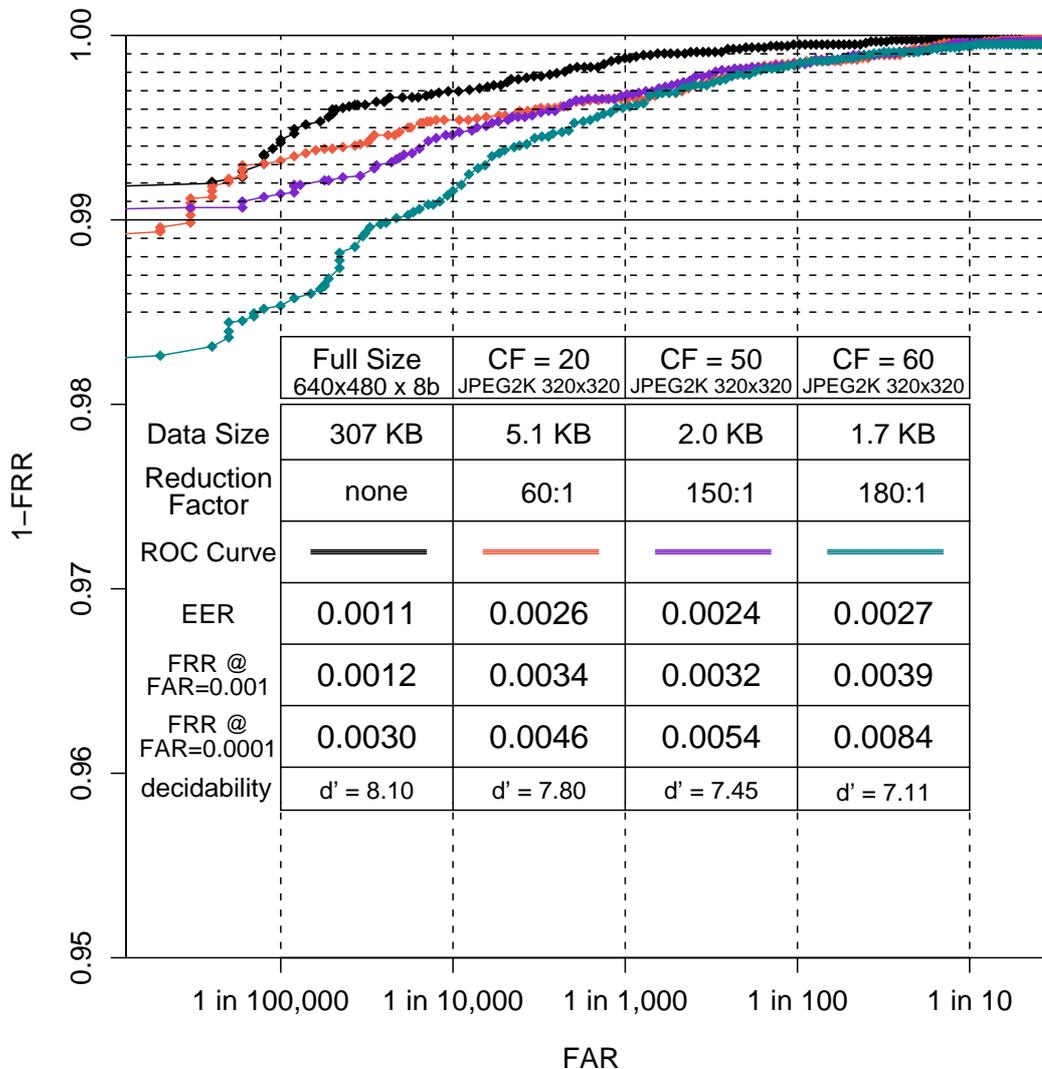


Figure 6: ROC curves and data size statistics showing iris recognition performance when the cropped and ROI-isolated images are compressed using JPEG2000 at various compression factors. Performance with file sizes of merely 2,000 bytes (CF = 50, purple curve) remains astonishingly unimpaired compared to baseline (black curve); but further compression exacts a high toll (blue-green curve).

## 5 Comparing the effects of the compression schemes

In this report 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 normalised Hamming Distance (HD), which is the fraction of

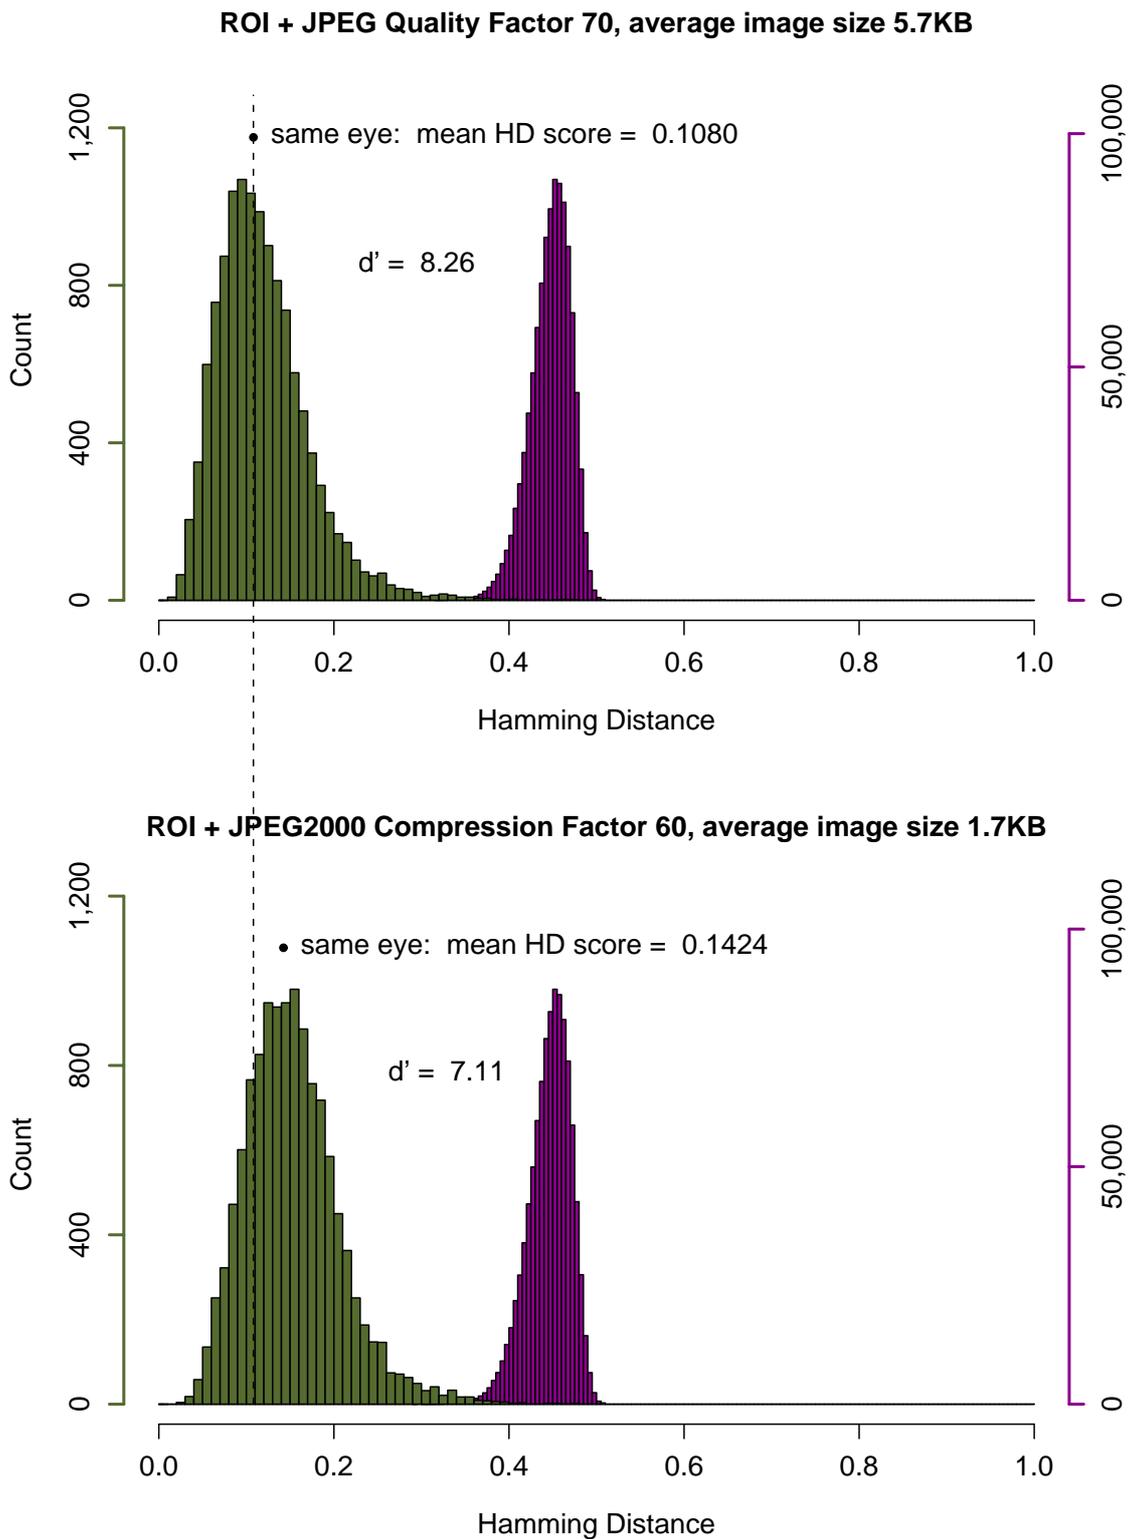


Figure 7: Distributions of Hamming Distance scores comparing same and different eyes in the NIST database, for two of the image compression schemes bracketing the range of schemes studied. Even in the most severe case (lower panel) using images compressed to only 1,700 bytes, the dual distributions have little overlap and so decisions about identity remain remarkably robust.

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 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 5,700 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 1,700 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 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 and infinite 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 non-vanishing theoretical models for them.

An alternative family of distance metrics, encompassing the Fisher ratio and Z-scores, define distance in terms of the difference between the means of the two distributions, normalised by some function of their standard deviations. One such is the  $d'$  metric of decidability in signal detection theory, defined as  $d' = |\mu_1 - \mu_2| / \sqrt{\frac{1}{2}(\sigma_1^2 + \sigma_2^2)}$ , where  $\mu_1$  and  $\mu_2$  are the means and  $\sigma_1$  and  $\sigma_2$  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, 4, 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.

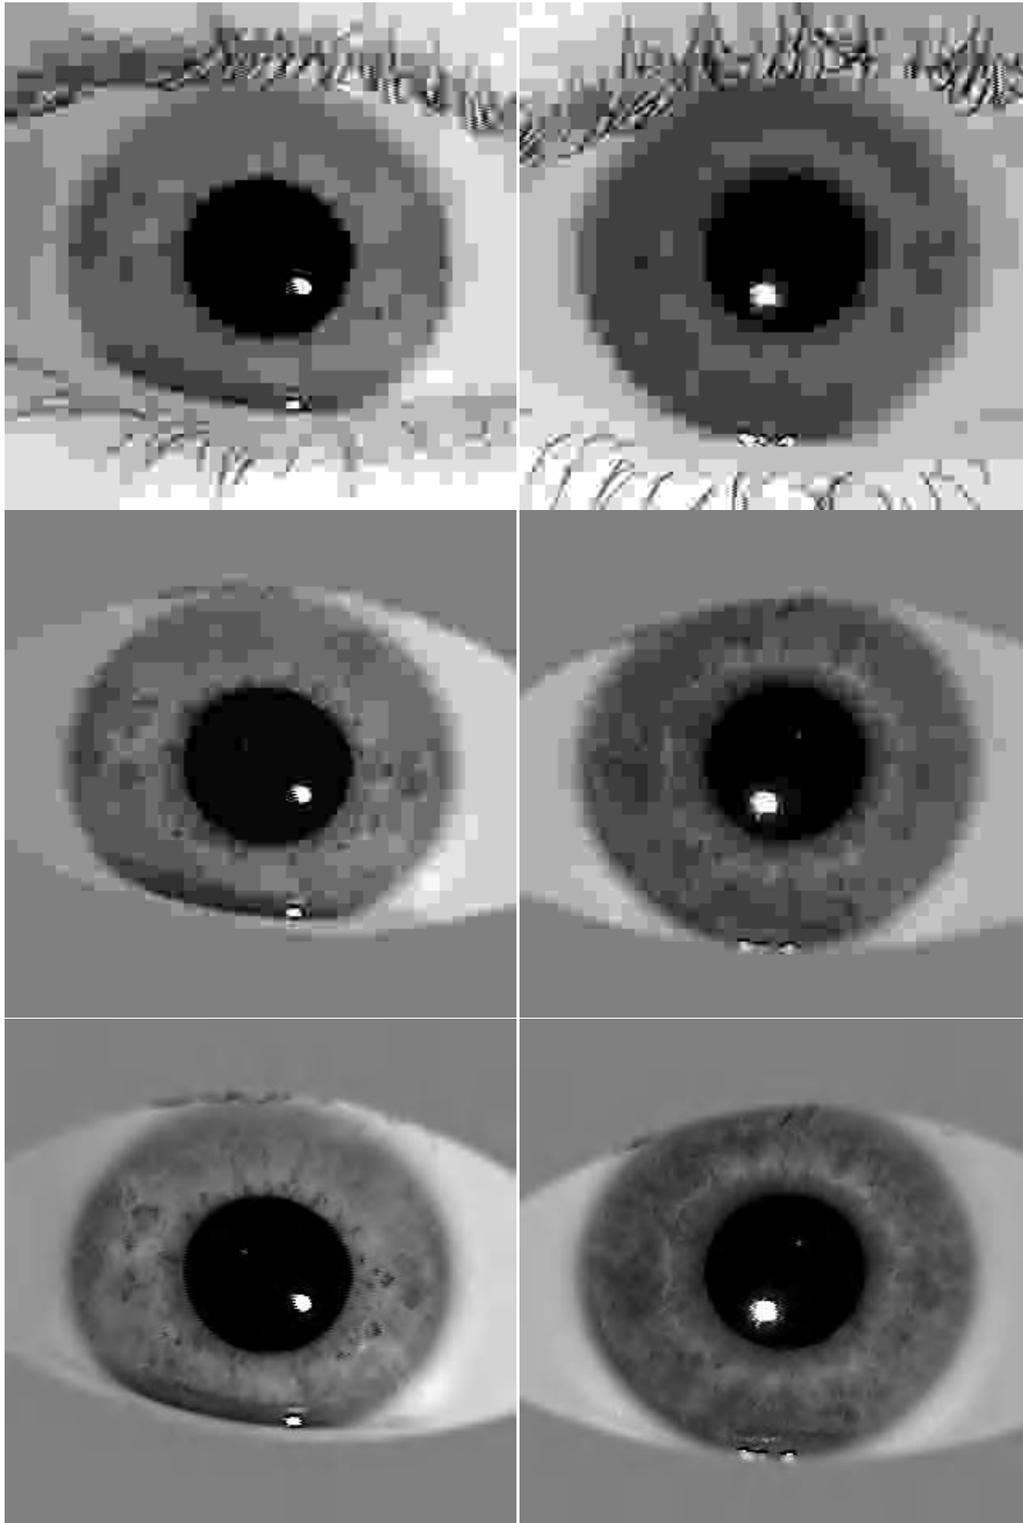


Figure 8: Visual comparison of the three schemes for iris image compression, for images all compressed to the same data size of 2,000 bytes. Left column is NIST 239230; right is NIST 239343. Top row: simple JPEG compression of the cropped ( $320 \times 320$ ) images. Middle row: JPEG compression of the 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 in the purple ROC curve ( $CF = 50$ ) in Fig. 6.

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.

Finally, it is interesting to compare visually some examples of the iris images after compression to a constant data size of 2,000 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 a cropped ( $320 \times 320$ ) image but without ROI isolation. Most of the 2,000 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 summarised by the purple ROC curve ( $CF = 50$ ) in Fig. 6.

## 6 Conclusions

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 originals, to just 2,000 bytes, remain perfectly 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 remains in rectilinear array form, 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 a final measure of how much impact each of the compression methods has on iris encoding, we compared the IrisCodes generated under each scheme to those generated for the corresponding original uncompressed images. The average HD (fraction of disagreeing bits) between such IrisCodes obtained before and after image compression is presented in Table 1 for each scheme and for each compression parameter, as interoperability scores. They indicate that only about 2% to 3% of the IrisCode bits change as a consequence of image compression even as severe as to 2,000 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 consequence indeed.

Strategy	Compression Parameter	Average Image Size	Interoperability Hamming Distance
Cropping (320 × 320) + JPEG Compression	QF = 70	<b>12.4 KB</b>	0.006
	QF = 30	<b>5.7 KB</b>	0.011
	QF = 20	<b>4.2 KB</b>	0.021
Cropping + ROI + JPEG Compression	QF = 70	<b>5.7 KB</b>	0.015
	QF = 30	<b>2.7 KB</b>	0.021
	QF = 20	<b>2.1 KB</b>	0.031
Cropping + ROI + JPEG2000 Compression	CF = 20	<b>5.1 KB</b>	0.018
	CF = 50	<b>2.0 KB</b>	0.027
	CF = 60	<b>1.7 KB</b>	0.035

Table 1: Summary of the compression schemes, resulting image file sizes, and the effect of compression on the computed IrisCodes, expressed as the fraction of bits that were changed from those computed for the original full-size images.

## References

- [1] C.E. Shannon, “A mathematical theory of communication,” *Bell Sys. Tech. Journal*, 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. on Pat. Anal. and Mach. Intell.*, vol. 15, no. 6, pp. 569–579, 1993.
- [4] R. Cappelli, D. Maio, and D. Maltoni, “Synthetic fingerprint-image generation,” *Proc. Int’l. Conf. on Pat. Recog. 15*, vol. 3, pp. 475–478, 2000.
- [5] J. Cui, Y. Wang, J. Huang, T. Tan, and Z. Sun, “An iris image synthesis method based on PCA and super-resolution,” *Proc. 17th Int’l. Conf. on Pat. Recog.*, vol. 4, pp. 471–474, 2004.
- [6] J. Zuo, N.A. Schmid, and X. Chen, “On generation and analysis of synthetic iris images,” *IEEE Trans. on Forensics and Information Security*, vol. 2, no. 1, pp. 77–90, 2007.

- [7] International Standard ISO/IEC 19794-6: “Information technology – Biometric data interchange formats – Part 6: Iris image data.” ISBN 0-580-46456-3, 2005.
- [8] Registered Traveler Interoperability Consortium (RTIC). “Technical Interoperability Specification for the Registered Traveler Program (Version 1.0 – Final),” 2006. Available at: <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. of SPIE (Applications of Digital Image Processing XIX)*, vol. 2847, 1996.
- [10] J. Daugman and C. Downing, “Epigenetic randomness, complexity, and singularity of human iris patterns,” *Proceedings of the Royal Society, B*, vol. 268, *Biological Sciences*, pp. 1737–1740, 2001.
- [11] J. Daugman, “How iris recognition works,” *IEEE Trans. Circuits and Systems for Video Technology*, vol. 14, no. 1, pp. 21–30, 2004.
- [12] S. Rakshit and D.M. Monro, “An evaluation of image sampling and compression for human iris recognition,” *IEEE Trans. on Forensics and Information Security*, in press, 2007.
- [13] National Institute of Standards and Technology, Iris Challenge Evaluation. Available from: <http://iris.nist.gov/ice/>
- [14] G.K. Wallace, “The JPEG still picture compression standard,” *Comm. of the ACM*, vol. 34, no. 4, pp. 30–44, 1991.
- [15] International Standard ISO/IEC 10918: “Information technology – Digital compression and coding of continuous-tone still images,” 1994. <http://www.jpeg.org>
- [16] A.P. Bradley and F.W.M. Stentiford, “JPEG2000 and region of interest coding,” *Digital Imaging Computing Techniques and Applications*, Melbourne, 2002. Available at: <http://www.ee.ucl.ac.uk/~fstentif/DICTA02.pdf>
- [17] International Standard ISO/IEC 15444-1: “Information technology – JPEG 2000 image coding system,” 2004. Available at: <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. Cons. Electr.*, vol. 46, no. 4, pp. 1103–1127, 2000. Available at: [http://jj2000.epfl.ch/jj\\_publications/papers/006.pdf](http://jj2000.epfl.ch/jj_publications/papers/006.pdf)
- [19] R.L.V. Hsu and P. Griffin, “JPEG region of interest compression for face recognition,” *IDENTIX Document No. RDNJ-04-0102*, 2005.