Custom Design of JPEG Quantisation Tables for Compressing Iris Polar Images to Improve Recognition Accuracy^{*}

Mario Konrad¹, Herbert Stögner¹, and Andreas Uhl^{1,2}

¹ School of Communication Engineering for IT, Carinthia Tech Institute, Austria

² Department of Computer Sciences, Salzburg University, Austria
uhl@cosy.sbg.ac.at

Abstract. Custom JPEG quantisation matrices are proposed to be used in the context of compressing iris polar images within iris recognition. These matrices are obtained by employing a Genetic algorithm for the corresponding optimisation. Superior matching results in iris recognition in terms of average Hamming distance and improved ROC are found as compared to the use of the default JPEG quantisation table.

1 Introduction

With the increasing usage of biometric systems the question arises naturally how to store and handle the acquired sensor data. In this context, the compression of these data may become imperative under certain circumstances due to the large amounts of data involved. Among other possibilities (e.g. like compressed template storage on IC cards), compression technology may be applied to sample data in two stages of the processing chain in classical biometric recognition for example: Transmission of sample data after sensor data acquisition and optional storage of (encrypted) reference data in template databases.

The distortions introduced by lossy compression artifacts usually interfere with subsequent feature extraction and may degrade the matching results. In particular, FRR or FNMR will increase (since features of the data of legitimate users are extracted less accurately from compressed data) which in turn affects user convenience and general acceptance of the biometric system. In extreme cases, even FAR or FMR might be affected.

In this work, we will focus on the lossy compression of iris polar images using the JPEG standard. We discuss the use of custom quantisation matrices in order to reflect the specific properties of iris imagery. We will apply a biometric iris recognition systems to the compressed sensor data to evaluate the effects of compression on recognition accuracy.

In Section 2, we will review and discuss the available literature on biometric sample data compression with focus on iris data storage. Section 3 is the main part of this work where we discuss properties of iris imagery and present several

^{*} This work has been supported by the Austrian Science Fund, project no. L554-N15.

variants of custom JPEG quantisation tables (designed in order to hopefully improve recognition accuracy). In section 4 we first describe the employed iris recognition system and the data this algorithm are applied to. Subsequently we discuss our experimental results with respect to the observed improvements of recognition accuracy. Finally, we describe the Genetic algorithm approach which has been employed for optimising the desired quantisation matrices. Section 5 concludes the paper.

2 Iris Image Compression

ISO/IEC 19794-6 allows iris image data to be stored in lossy manner in the JPEG and JPEG2000 formats. Two types of iris image data are considered: rectilinear images (i.e. images of the entire eye) and polar images (which are basically the result of iris detection and segmentation), the latter much smaller in terms of storage requirement (e.g. 2kB vs. 25-30kB for rectilinear images).

In previous work on compressing iris imagery, rectangular [1,2,3,4] as well as polar [5] iris sample data has been considered. With respect to employed compression technology, we find JPEG [2,3,4], JPEG2000 [1,5,2,3,4], and other general purpose compression techniques [3,4] being investigated. Superior compression performance of JPEG2000 over JPEG is seen especially for low bitrates, however, for high and medium quality JPEG is still an option to consider. While the data formats specified by the ISO/IEC 19794 standard are fixed at present state, their customised use tailored to a specific target modality and the corresponding impact on recognition accuracy as compared to the default settings has not been investigated. In the subsequent study we apply JPEG as covered by ISO/IEC 19794-6 to polar iris images and propose to use custom quantisation tables (Q-tables) adapted to properties of iris imagery.

In some application settings, the requirement for compression technology is caused by low power (mobile) sample acquisition devices which are too weak to conduct feature extraction on board and therefore need to transmit sample data to a remote feature extraction (and matching) module. In this context, it of course makes more sense to apply JPEG instead of JPEG2000 due to its much lower computational demand. In addition to this, applying compression to polar iris images minimises the amount of data to be transmitted (since polar images are smaller by several orders of magnitude as compared to rectangular iris images even without compression applied). Also, this strategy avoids the iris detection process being fooled by compression artifacts as it may be the case when the iris needs to be detected in compressed rectangular iris images. A drawback of the approach relying on polar images is that the acquisition device needs to perform iris detection and the generation of the iris texture patch (i.e. the polar iris image) which involves data interpolation or extrapolation. In any case, the bandwidth required for transmission of sample data is minimised by employing compressed polar iris data.

In [6] compression algorithms tuned for application in the pattern recognition context are proposed, which are based on the modification of the standard compression algorithms: This is done by emphasising middle and high frequencies

and discarding low frequencies (the standard JPEG Q-table is rotated by 180 degrees). JPEG Q-table optimisation has already been considered in biometrics – [7] employ a rate/distortion criterion in the context of face recognition and achieve superior recognition performance as compared to the standard matrix.

3 Custom JPEG Quantisation

The JPEG still image compression standard [8] allows to use custom Q-tables in case image material with special properties is subject to compression. These tables are signalled in the header information. The default Q-tables (see Table 1) have been designed with respect to psychovisual optimality employing large scale experimentation involving a high number of test subjects. There are two reasons which suggest to use different Q-tables as the default configuration: First, iris imagery might have different properties as compared to common arbitrary images, and second, a pleasant viewing experience as being the aim in designing the default tables, might not deliver optimal matching results in the context of biometric recognition (e.g. sharp edges required for exact matching could appear appealing to human observers).

Therefore, as a first stage, we have investigated iris imagery in more detail. 8x8 pixel image blocks have been subjected to DCT transform and the resulting coefficients are averaged for a large number of blocks (i.e. 2000 and 525 blocks for the two types of imagery, respectively). As a first class of blocks, we have used arbitrary images and blocks are extracted randomly. The second class of blocks is taken from polar iris images. Fig. 1 displays the result of both classes where the DC and the largest AC coefficient are set to white, zero is set to black and the remaining values are scaled in between (note that the logarithm is applied to the magnitude of all coefficients before this scaling operation).

The arbitrary blocks (Fig. 1.a) show the typical expected behaviour with decreasing coefficient magnitude for increasing frequency and symmetry with

Fig. 1. Averaged 8x8 DCT blocks

respect to the coordinate axes. Fig. 1.b reveals that in polar iris images there is more energy in the higher frequencies in horizontal direction as compared to vertical direction. This is to be expected since luminance fluctuations in iris texture are more pronounced in radial direction as compared to perpendicular direction.

While we may exploit the direction bias of iris texture in compression directly, we additionally conjecture that the highest and medium frequencies might not be required for the matching stage due to the coarse quantisation used for template generation while at least medium frequencies are required for pleasant viewing. Table 1 displays the Q-tables used in our experiments.

Table 1. JPEG Quantisation tables: STQ, Qtable22, Qtable24 (first line), QTOptk05, and QTOptk10 (second line)

16 11 10 16 24 40 51 61	10 10 76 255 255 255 255 255	16 11 10 16 255 255 255 255
12 12 14 19 26 58 60 55	85 112 255 255 255 255 255 255	12 12 14 255 255 255 255 255
14 13 16 24 40 57 69 56	151 255 255 255 255 255 255 255	255 255 255 255 255 255 255 255
14 17 22 29 51 87 80 62	255 255 255 255 255 255 255 255	255 255 255 255 255 255 255 255
18 22 37 56 68 109 103 77	255 255 255 255 255 255 255 255	255 255 255 255 255 255 255 255
24 35 55 64 81 104 113 92	255 255 255 255 255 255 255 255	255 255 255 255 255 255 255 255
49 64 78 87 103 121 120 101	255 255 255 255 255 255 255 255	255 255 255 255 255 255 255 255
72 92 95 98 112 100 103 99	255 255 255 255 255 255 255 255	255 255 255 255 255 255 255 255
16 11 10 16 24 246 255 255	15 6 17 19 255 255 255 255	
12 12 14 29 26 255 255 250	5 15 13 255 255 255 250 255	
14 13 16 24 255 255 255 224	255 255 247 255 248 255 255 255	
14 17 22 255 255 255 242 255	255 250 248 255 255 250 255 255	
18 255 255 255 255 255 255 255	255 222 237 255 251 255 255 250	
24 247 255 255 255 255 255 255	255 252 251 250 220 249 229 232	
255 255 255 255 255 255 255 255	254 246 255 251 255 255 255 248	
255 255 255 241 255 241 255 244	255 255 247 252 255 255 248 255	

The first matrix shows the standard case (STQ) where the entries exhibit a steady increase from low frequencies to high frequencies following the well-known zig-zag pattern [8] (which results in more severe quantisation applied to middle and high frequencies).

Qtable22 and Qtable24 have been obtained by large scale trial and error experimentation, setting a large amount of entries to 255 (which causes the corresponding coefficients to be divided by 255 and results in most of them being quantised to zero). Both matrices are asymmetric in the sense that they “protect” more coefficients in horizontal direction (which have been shown to carry more energy as their vertical counterparts in Fig. 1.b), Qtable24 is more pronounced in this respect and retains the values of STQ at the positions not set to 255. The rationale behind the selection of these matrices is to investigate the importance of medium frequency information in the iris recognition process (high frequency information is assumed to be not useful in any case) and to reflect the specific properties of polar iris images.

QTOptk05 and QTOptk10 have been found using the Genetic optimisation approach as described in Section 4.3 using Qtable22 and Qtable24 as individuals of the initial population in addition to randomly generated tables. These tables have been specifically optimised for application with compression rates 5 and 10, respectively.

4 Experimental Study

4.1 Setting and Methods

Iris Recognition System. The employed iris recognition system is Libor Masek's Matlab implementation¹ of a 1-D version of the Daugman iris recognition algorithm. First, this algorithm segments the eye image into the iris and the remainder of the image. Iris image texture is mapped to polar coordinates resulting in a rectangular patch which is denoted "polar image". In case compression is required, it is applied to the polar image in this stage of the procedure. After extracting the features of the iris (which are strongly quantised phase responses of complex 1-D Gabor filters in this case), considering translation, rotations, and disturbed regions in the iris (a noise mask is generated), the algorithm outputs the similarity score by giving the Hamming distance between two extracted templates.

Sample Data. For all our experiments we considered 320x280 pixel images with 8-bit grayscale information per pixel from the CASIA² 1.0 iris image database. Note that the fact that the pupil area has been manipulated in these data [9] does not affect our results since we restrict compression to the iris texture area only by compressing polar iris images. We applied the experimental calculations on the images of 50 persons using 3 - 4 images for each eye (i.e. 334 images).

Fig. 2 shows examples of iris templates extracted from uncompressed (first line) and JPEG compressed iris texture patches of one person. Second and third patch (second line) are compressed with rate 10 using STQ and QTOptk10, respectively. By analogy, fourth and fifth patch (third line) are compressed with rate 15 using STQ and Qtable22, respectively.

Note that the images have been scaled in y-direction for proper display, the original dimension is 240×20 pixels.

In the iris texture data (not shown), compression artifacts are clearly visible for both rates displayed, however, the STQ compressed variants are visually more close to the original and seem to have conserved the texture details much better. However, when computing the Hamming distance between both variants having applied compression rate 15 with the uncompressed second image of the same eye in the database, we result in 0.327 for STQ but only 0.317 for Qtable22. Obviously, the Hamming distance between templates does not reflect the visual appearance at all. The smoothing achieved by compression seems to play an important role indeed. More fine grained differences seem to get introduced by the STQ quantisation, while the other two matrices tend to produce rather smooth templates as well.

Compression can be used in various stages of the recognition/matching process. Either the stored reference data may be in compressed form, or the sample data acquired for verification may be compressed (e.g. for efficient transfer), or

¹ <http://www.csse.uwa.edu.au/~pk/studentprojects/libor/sourcecode.html>

² <http://www.sinobiometrics.com>

Fig. 2. Comparison of iris templates extracted from uncompressed (top line)/ compressed polar iris images with compression rates 10 (left column) and 15 (right column)

both. Therefore, we use two settings in our experiments: either both images are compressed and matched against each other or only one image is compressed in the matching stage. For investigating correct matches (matches from legitimate users enrolled in the database), we rely on more than 12000 generated images (i.e. 826 images for genuine user matches per compression rate; considering the 15 different compression rates applied (rates 2 - 16), we finally result in 12390 overall images considered). This is only true in the scenario with only 1 compressed image, for 2 compressed images this number is half-ed due to symmetry reasons. For investigating matches between different persons (imposter matches), far more data is available of course (109304 imposter matches are considered for a single rate).

4.2 Experimental Results

First, we investigate the impact of compression on the matching score (i.e. obtained Hamming distance (HD)). Fig. 3 shows the plots of the HD after applying the iris recognition algorithm if iris polar images have been JPEG compressed in the case of genuine user matches. The x-axis shows the compression rates, whereas the y-axis shows the averaged Hamming distance. For reference, we have included the average HD for the case of uncompressed images as horizontal dashed line with circles (as it is the case for all subsequent plots).

The mean value of the HD in the uncompressed case is approximately 0.3. First we consider the standard Q-table (labelled STQ). For increasing compression rate the average HD increases steadily and crosses the suggested matching threshold of 0.34 at compression rates 12 for both cases (one or two images compressed, respectively). Note that the reported numbers refer to averaged HD values which implies the occurrence of a significant number of false negative matches at this compression rate.

Fig. 3. Impact of varying compression rate on HD of genuine users' matches

Concerning the one compressed image scenario (Fig. 3.a), STQ is beaten by QTOptk05 (rates around 6) and Qtable24/Qtable22 (rates 14 and higher), but only by a very small amount.

The situation is different when regarding the two compressed images scenario (Fig. 3.b). QTOptk05 is clearly better than STQ between rate 4 and 8 (and even beats the uncompressed case between rate 4 and 7). Qtable24 and QTOptk10 are better than STQ for rates higher than 7 and also beat the uncompressed case between rate 8 and 11. Finally, Qtable22 beats STQ for rates higher than 10 and is also superior to the uncompressed case for rates 13 and higher.

Next, we focus on rate/distortion performance in terms of PSNR. Figure 4.a shows the averaged rate distortion comparison of JPEG applied to all iris images for the five Q-tables considered. As it is the case for all subsequent plots, the solid graph depicted with crosses shows the results of the standard matrix (STQ). Some interesting results may be observed. First, Qtable24 behaves similarly to QTOptk10 and both exhibit PSNR values clearly above STQ for compression rates larger than 9 – up to 2dB difference may be observed. QTOptk05 is slightly above STQ between rate 5 and rate 8, but the improvement seen is only up to 0.4 dB. Qtable22 shows very fluctuating behaviour for low compression rates, but significantly outperforms STQ for rates larger than 11, up to 2.5dB improvement is found especially for higher rates.

Interestingly, we find that PSNR behaviour is highly dependent on the rate considered and all investigated quantisation matrices are able to outperform STQ in a certain range considerably.

These results indicate that PSNR is indeed a good predictor for matching performance with two compressed iris images in terms of average Hamming distance, but **NOT** in the case of only one image being compressed. The claim that compression up to a rate of 16 even improves the matching scores of not compressed images [5] can be supported at least for the 2 compressed images case for certain “better” Q-tables in distinct ranges of compression rate. This fact is remarkable and may be explained by the fact that compression acts as a kind of low pass filter resulting in denoised and smoothened images which can be matched better as compared to the original noisy counterparts.

Fig. 4. Rate/distortion and ROC performance

Fig. 5. ROC at different rates

In order to consider the hidden statistical outliers in the comparisons and to use a quantity often employed in the assessment of biometric system performance, we focus on the receiver operating characteristic (ROC) by computing and plotting the false rejection rate (FRR) against the false acceptance rate (FAR) for different compression rates.

Figs. 4.b to 5 compare the ROC of the different Q-tables for compression rates 5, 10, and 15 (it does not seem to be realistic to operate the iris recognition system at a higher compression rate due to the low visual quality of the images – see Fig. 2). We focus on the two compressed image scenario since the effects observed are identical to the one compressed image case but are seen in a more pronounced manner.

For compression rate 5 (see Fig. 4.b), the proposed QTOptk05 is able to improve the uncompressed and STQ ROC at $FAR < 0.1$ and $FRR > 0.02$. Note also that STQ does hardly outperform the uncompressed case, while this is observed for QTOptk05.

In the case of compression rate 10 (see Fig. 5.a), the situation changes drastically. Now, Qtable24 and QTOptk10 improve over the uncompressed case at $FAR < 0.1$ and $FRR > 0.015$, while QTOptk05 now performs almost equally to STQ and clearly inferior to the uncompressed case.

Finally, when turning to compression rate 15 the situation is again different (Fig. 5.b): now the uncompressed ROC is better as all compressed variants. However, Qtable22 is rather close to the corresponding curve. When comparing ROC to the STQ case, we clearly observe that the customised tables significantly improve over STQ in the entire range displayed in the plots.

There is one more interesting thing to note: STQ performs worst of all investigated matrices. With this rather high compression rate, Qtable22 offers the possibility to actually use the recognition algorithm whereas for STQ the ROC behaviour is actually too poor to be applied in any practical setting.

4.3 Genetic Algorithm Optimisation Approach

We have employed a Genetic algorithm (GA) to generate the two matrices QTOptk05 and QTOptk10 as follows. The Q-table entries (we have restricted the values to be integers from the interval $[0, 255]$) constitute the genes of each individual, where an individual represents a distinct Q-table. The population size is set to 10 individuals and we limit the number of computed generations to 40 due to reasons of computational demand. Additionally, the optimisation is stopped if no improvement in terms of best individual fitness function is found for 10 generations. In each generation, the two best individuals are selected to be the “elite” and are kept for the subsequent generation. Six individuals are selected for crossover, while two individuals of the subsequent generation are created by mutation.

As the cost function to be evaluated for determining which individuals are to be kept, we compute the sum of the following items for a fixed compression rate: averaged genuine users Hamming distance, average of FAR over a selected

Fig. 6. Cost function development in the GA (compression rate 10)

number of thresholds, and average of FRR over the same set of thresholds. This costfunction has to be minimised of course.

Figure 6 shows the development of the cost function values for two cases: the mean of the fitness computed over each generation and the best fitness value in each generation. The mean fitness function value is not further improved after an initial “jump”, while the best value is improved in several steps until saturation is reached after 30 generations and the GA stops. Note that we have used well performing tables like Qtable22 and Qtable24 as parts of the initial population in addition to randomly generated individuals.

We speculate that a higher mutation rate and a more disruptive crossover strategy might lead to even better results and will conduct further experiments in this direction.

5 Conclusion and Future Work

We have found that custom designed quantisation tables in JPEG can improve matching results in terms of average HD and ROC behaviour significantly as compared to the default tables. This effect is more pronounced for higher compression rates and for the scenario where both images involved in matching are compressed. Moreover it has turned out that these custom matrices need to be optimised with respect to a specific target bitrate – significant improvements are only found within the bitrate the table has been optimized for.

In future work we will consider additional alternative iris recognition algorithms in order to identify possible interference between compression technique and iris recognition system. Furthermore we will further optimise GA parameters in order to determine the ideal configuration.

Acknowledgements

Most of the work described in this paper has been done in the scope of the “Project I Lab” in the master program on “Communication Engineering for IT” at Carinthia Tech Institute.

References

1. Ives, R.W., Broussard, R.P., Kennell, L.R., Soldan, D.L.: Effects of image compression on iris recognition system performance. *Journal of Electronic Imaging* 17, 011015 (2008)
2. Daugman, J., Downing, C.: Effect of severe image compression on iris recognition performance. *IEEE Transactions on Information Forensics and Security* 3(1), 52–61 (2008)
3. Matschitsch, S., Tschinder, M., Uhl, A.: Comparison of compression algorithms’ impact on iris recognition accuracy. In: Lee, S.-W., Li, S.Z. (eds.) *ICB 2007*. LNCS, vol. 4642, pp. 232–241. Springer, Heidelberg (2007)

4. Jenisch, S., Lukesch, S., Uhl, A.: Comparison of compression algorithms' impact on iris recognition accuracy II: revisiting JPEG. In: Proceedings of SPIE, Security, Forensics, Steganography, and Watermarking of Multimedia Contents X, San Jose, CA, USA, vol. 6819, p. 68190M (January 2008)
5. Rakshit, S., Monro, D.: Effects of sampling and compression on human iris verification. In: Proceedings of the IEEE International Conference on Acustics, Speech, and Signal Processing (ICASSP 2006), Toulouse, France, pp. II-337-II-340 (2006)
6. Chen, M., Zhang, S., Karim, M.: Modification of standard image compression methods for correlation-based pattern recognition. *Optical Engineering* 43(8), 1723-1730 (2004)
7. Jeong, G.M., Kim, C., Ahn, H.S., Ahn, B.J.: JPEG quantization table design for face images and its application to face recognition. *IEICE Transactions on Fundamentals of Electronics, Communications and Computer Science* E69-A(11), 2990-2993 (2006)
8. Pennebaker, W., Mitchell, J.: JPEG – Still image compression standard. Van Nostrand Reinhold, New York (1993)
9. Philips, P., Bowyer, K., Flynn, P.: Comments on the CASIA version 1.0 iris data set. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 29(10), 1869-1870 (2007)