SELECTIVE JPEG2000 ENCRYPTION OF IRIS DATA: PROTECTING SAMPLE DATA VS. NORMALISED TEXTURE

Martin Rieger, Jutta Hämmerle-Uhl, and Andreas Uhl

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

ABSTRACT

Biometric system security requires cryptographic protection of sample data under certain circumstances. We assess low complexity selective encryption schemes applied to JPEG2000 compressed iris data by conducting iris recognition on the selectively encrypted data. This paper specifically compares the effects of a recently proposed approach, i.e. applying selective encryption to normalised texture data, to encrypting classical sample data. We assess achieved protection level as well as computational cost of the considered schemes, and particularly highlight the role of segmentation in obtaining surprising results.

Index Terms— Iris sample protection, selective encryption, JPEG2000, iris recognition

1. INTRODUCTION

The International Organisation for Standardisation (ISO) specifies biometric data to be also recorded and stored in (raw) image form (ISO/IEC FDIS 19794), not only in extracted templates (e.g. minutiae-lists or iris-codes). The certainly most relevant standard for compressing such data is JPEG2000, suggested for (lossy) compression of iris sample images in the ISO/IEC 19794 standard suite on Biometric Data Interchange Formats. As these data are highly privacy sensitive, cryptographic protection for (long-term) storage and additionally e.g. for transmission between sensor and authentication module is required.

As this application context does not require protected matching (as requested from template protection schemes), it allows the employment of classical cryptographic techniques.

In this paper we investigate lightweight encryption schemes for JPEG2000 compressed iris sample data, suited for mobile and/or low-power environments, based on selective bitstream protection applied to either sample data or normalised iris texture. While the latter approach suggested recently [1] is appealing due to the low amount of data to be compressed and protected, it also exhibits certain disadvantages which will be discussed and evaluated.

This work has been partially supported by the Austrian Science Fund, project no. 27776.

Section 2 introduces principles of encrypting JPEG2000 data and specifically describes the approaches tailored for iris data as proposed in this paper. Section 3 describes conducted experiments, where we specifically assess the security of the proposed encryption schemes by applying four different iris recognition schemes to the (attacked) encrypted data. The role of iris segmentation on protected data is particularly highlighted. Section 4 presents the conclusions of this paper.

2. EFFICIENT ENCRYPTION OF IRIS DATA

2.1. Iris Sample Data Types

The iris recognition processing chain typically consist of several stages, the first of which is iris localisation also termed "iris segmentation" where the pupillar and limbic boundaries of iris texture are determined.

In the second step the localised iris is normalised. The reasons for this are differences in image acquisition, like the varying size of irises caused by changes of the camera-to-eye distance. The area between the two boundary-curves is mapped into a rectangle texture with fixed size for compensating such deformations using a coordinate-transform from Cartesian-coordinates to polar-coordinates also denoted as "rubber sheet-transform" (see Fig. 2.a). The final pre-processing step enhances contrast and compensates for illumination variations by applying e.g. CLAHE to the normalised texture.

The experiments are done on the CASIA V3 Interval dataset. The original samples have a resolution of 320×280 pixels with 8bpp grayscale (NIR data), while the normalised iris texture derived using USITv2[1] (University of Salzburg Iris Toolkit v2.0.x [2, 3]) has a resolution of 512×64 pixels with identical bitdepth, thus the pixel count is reduced by a factor of 2.73 when considering normalised iris texture. Note that these two types of iris data correspond to standardised iris images (IREX records) as defined by the NIST Iris Exchange (IREX I `http://iris.nist.gov/irex/`) program. In particular, original sample data corresponds to IREX record kind 1 or 3, while the normalised texture corresponds to record kind 16 (which has been later abandoned by NIST).

[1] `http://www.wavelab.at/sources/USIT/`

The observation of reduced data rate for normalised iris textures motivates the employment of this sample data representation when low computational cost is aimed for.

Table 1 compares the filesize of the two different data types after lossless JPEG2000 compression, i.e. comparing the amount of data subjected to encryption in case of full protection is being applied.

Table 1: Filesize in byte after JPEG2000 lossless compression (CASIA V3 Interval).

Data	$\varnothing$	σ	range
original sample	42501.80	3874.06	[27402,51998]
normalised texture	15471.84	1119.22	[10408,19803]

We observe that the relation between original sample compressed file size and normalised texture compressed file size is preserved from the case of looking at image dimensions only, and also the file size variability is significantly lower for normalised textures (which is beneficial for worst case planning). Thus, encrypting normalised textures looks like a pretty good idea as proposed in [1].

2.2. Selective JPEG2000 Encryption Approaches and Security Assessment

For JPEG2000, [4] provides a comprehensive survey of encryption schemes. In our target application context, only bitstream oriented techniques and format compliant ones (to enable security assessment of encrypted data) are appropriate. We apply a corresponding scheme introduced in the context of JPSEC [5]).

In a series of papers (i.e. [6, 7, 8]) we have defined and analysed different ways how to apply encryption to different parts of a fingerprint-image JPEG2000 codestream. From these techniques, we adopt "Absolute Encryption" (encryption is applied to one single chunk of data right at the *begin* of the codestream [6]) as this has proven to be the most sensible approach due to the progressive nature (most important information is concentrated at the start of the bitstream) and embeddedness of the JPEG2000 bitstream structure [7, 8]. This approach is thus termed "begin" in some plots in this work.

When assessing the security of format compliantly encrypted visual data, the data can simply be decoded with the encrypted parts (called "direct decoding"). The encrypted parts introduce noise-type distortions into the data which kind of overlay the visual information still present in the data (see Figs. 1 and 2 left side of each pair). An informed attacker can do much better by removing the encrypted parts before decoding and replacing them by suited data minimising error metrics. This can be done most efficiently using codec specific error concealment tools, which treat encrypted data like any type of bitstream error ("error concealment attack"). As these tools have been developed by JPEG2000 specialists it is questionable if an attacker might do any better. Thus, any serious security analysis needs to consider encrypted imagery being attacked using this error concealment approach at least. The JJ2000 version used in the experiments includes the patches and enhancements to JPEG2000 error concealment provided by [9, 7], and results obtained by error concealment are denoted by "rep" (for replacement) in the result Figs. 6 and 7.

In Figs. 1 and 2 we provide visual examples for encrypted iris sample data and normalised texture, respectively. The left-sided image of each pair is directly decoded, while for the right-sided image error-concealment decoding is done.

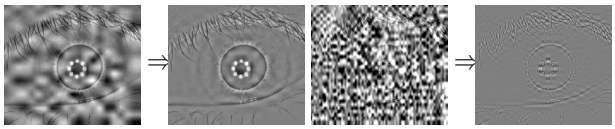

(a) 1% Absolute encryption (b) 5% Absolute encryption

Fig. 1: Comparison of encrypted iris sample images (direct decoding) to error-concealment decoding.

When comparing the obtained data after direct and error-concealment decoding, respectively, it gets immediately clear that security judgement based on direct reconstruction severely overestimates security (as the texture data still present is hidden by image noise). After error-concealment decoding most of the encryption noise is removed and in many cases iris-texture related structures are exhibited.

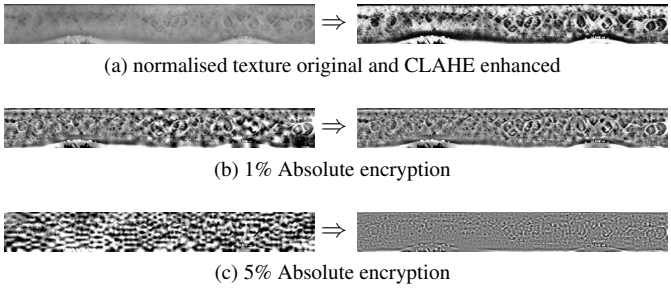

(a) normalised texture original and CLAHE enhanced

(b) 1% Absolute encryption

(c) 5% Absolute encryption

Fig. 2: Normalised texture: Original (a) and comparison of encrypted normalised textures (direct decoding) with error-concealment decoding (b) & (c).

In this work, actual security assessment is done by applying iris recognition schemes to the protected data (either after direct reconstruction or after having applied error-concealment decoding) to verify if the protection is sufficiently strong to prevent the use of the encrypted iris data in an automated recognition context.

3. EXPERIMENTS

3.1. Experimental Settings

All experiments are based on images taken from the CASIA V3 Interval iris image dataset consisting of 2647 NIR images from 395 different classes. Images, no matter if sample data or normalised texture, are compressed into lossless

JPEG2000 format using JJ2000 in resolution progressive ordering, using a single quality layer. The JPEG2000 bitstreams are encrypted by varying the amount of encrypted data and starting right from the bitstream start. Subsequently, data are either directly decoded or decoded with enabled error concealment with the JJ2000 variant mentioned [9]. *Segmentation* is performed using a method based on contrast-adjusted Hough transform (`caht`) extracting circular boundary curves proposed by [2].

For *feature extraction and matching*, we employ four techniques very different wrt. the dominant analysis orientation and the extraction domain considered, respectively (as different techniques might react in distinct manner to encryption artefacts): "**Ma**" [10], "**Masek**" [11], "**Ko**" [12], and "**Monro**" [13]. For a detailed description of our implementation of preprocessing, feature extraction, and matching see [2, 3]. All implementations are available in USITv2.

Iris recognition is then conducted in three variants: (i) directly applied to encrypted normalised texture (as suggested as low-cost encryption approach in [1]), (ii) applied to normalised texture which is generated from the encrypted sample by applying segmentation and normalisation to it (termed "classical segmentation") and (iii) applied to normalised texture generated from the encrypted sample when providing segmentation parameters computed on unprotected imagery and applied to the encrypted samples, and subsequent normalisation (termed "guided segmentation"). We consider both (ii) and (iii) to isolate the effect of segmentation applied to encrypted imagery as (iii) is not influenced by segmentation errors on encrypted data. While (i) and (iii) look pretty similar at first sight, the difference is that normalisation (including interpolation) is applied to encrypted data (and thus artefacts should propagate) in (iii) while in (i), encryption is applied to the already normalised data.

Error analysis is conducted by ROC analysis of iris recognition done in verification mode using the FVC2004 protocols [14], presenting EER always matching plaintext gallery images to encrypted probe images. Obviously, higher EER corresponds to better data protection.

3.2. Experimental Results

In the three following plots, we compare a direct reconstruction (left plot) to applying error-concealment during decoding the encrypted data (right plot).

Fig. 3, showing the results of applying recognition to the encrypted normalised texture, drastically illustrates that there is indeed a significant difference in the security assessment between considering direct reconstruction and error-concealment (an informed attacker in the latter case). The differences in recognition performance on plaintext data (Ma and Masek are clearly superior to Monro and Ko, respectively, see the values at 0% encryption on the x-axis) are clearly reflected also in the EER results on selectively protected data when error-concealment decoding is done (right plot). Results indicate that decent protection under Ma recognition is achieved after having encrypted at least 25% of the packet data. Also with Masek recognition, encrypting 10% of the packet data does not yet lead to the desired protection level (EER is still down to 25%). On the other hand, under Monro and Ko recognition, EER is up to 42% - 44% after having encrypted 4% of the packet data only.

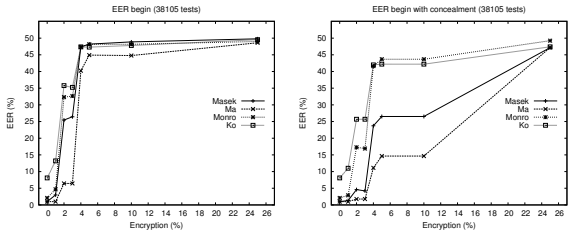

Fig. 3: Absolute encryption of normalised texture: EER using direct reconstruction and with error-concealment.

When comparing these results to those achieved after direct reconstruction (left plot), we see a very different picture. Results suggest that for all recognition types an encryption of 5% of the packet data is sufficient to result in more than 45% EER (more than 40% for Ma), thus indicating sufficiently secured data. These results drastically underline the importance of considering informed attackers to prevent an over-optimistic security assessment is resulting from this plot.

A very different picture is obtained when considering the case of applying recognition to encrypting sample data, i.e. when segmentation is applied to the encrypted data before normalisation (see Fig. 4). For both direct reconstruction and error-concealment application, respectively, EER is up to 50% already when encrypting 1% of the bitstream packet data only. This is a strong and somewhat surprising result.

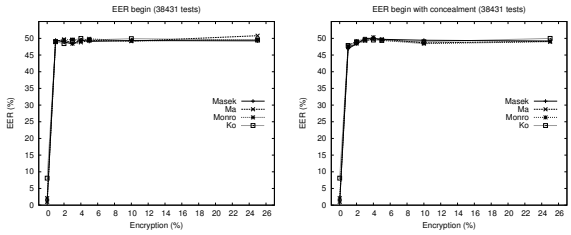

Fig. 4: Absolute encryption of samples with classical segmentation: EER using direct reconstruction and with error-concealment.

Note that encrypting this small amount of sample data roughly corresponds to encrypting 2.75% of normalised textures. When relating this to the normalised texture encryption results in Fig. 3, it is evident that encrypting sample data is the much better choice (as decent protection is achieved by encrypting clearly more than 10% of these data).

The guided segmentation scenario is meant to shed light on this result by excluding segmentation as a cause for introduced errors (as the segmentation result obtained on clear data

is applied to the encrypted sample data to rule out segmentation errors caused by encryption). The results are shown in Fig. 5. It is obvious, that the achieved protection is much weaker as compared to the previous case. Consequently, the high error rates observed in the classical segmentation setting are due to segmentation errors or probably failures.

When comparing encryption of normalised textures and encryption of samples with guided segmentation we notice better protection for the former case as high error rates for Ko and Monro are seen for a lower percentage of data encrypted (in the error-concealment case), and Ma in direct reconstruction achieves high EER at a later stage as compared to encrypting normalised textures.

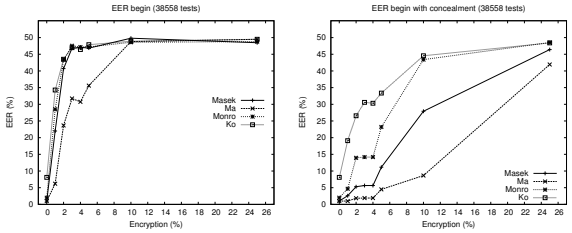

Fig. 5: Absolute encryption of samples with guided segmentation: EER using direct reconstruction and with error-concealment.

Also Masek and Ma (when considered under error-concealment reconstruction) achieve EER of 25% and 15% already when encrypting 5% of the bitstream in the case of encrypting normalised textures, while at this encryption amount the guided segmentation approach on encrypted samples delivers much lower EER. Note that this is specifically remarkable, as the absolute amount of data encrypted is a factor of over 2.5 higher for the encryption of sample data. Obviously, the application of normalisation including interpolation to encrypted data as done in the guided segmentation approach weakens the protection.

Finally we want to answer the question why the application of segmentation to encrypted samples leads to the observed superior protection properties. Fig. 6 displays the deviation of the centers of the detected boundary circles as well as the observed difference in radii magnitude for increasing encryption strength.

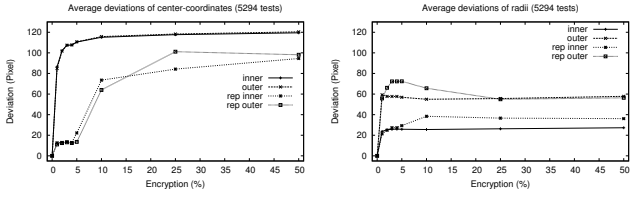

Fig. 6: Averaged deviation of center location and radii in direct reconstruction and error-concealment application.

For direct reconstruction we observe a complete segmentation disaster in terms of center displacement, but also for error-concealment application more than 10 pixels displacement are seen already when encrypting 1% of packet data, rising substantially for more than 5% encryption. For the radii, we observe the interesting effects that (i) the inner (pupil boundary) radius is more stable and (ii) that error-concealment reconstruction is worse compared to direct reconstruction. (i) does not come as a big surprise as the contrast of the pupil boundary is higher compared to the limbic one and (ii) might be explained by the smoothing involved in error concealment that probably makes boundary detection more difficult.

A strange phenomenon is the saturation and even decrease of the radii deviation for increasing encryption strength. However this has to be taken with care as shown in Fig. 7 (right), which illustrates the growing number of complete failures in circle detection for increasing encryption strength, in particular under error concealment. For direct reconstructions, the failure rate is much lower. This of course blurs the results on radii deviation and center displacement.

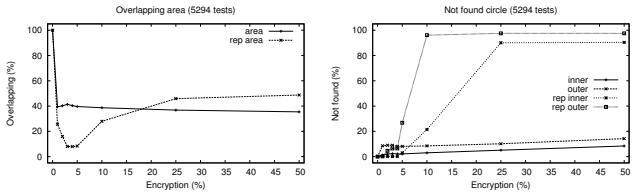

Fig. 7: Overlapping area and boundary detection failures.

Finally the left plot in Fig. 7 shows the percentage of pixels shared between iris texture area when segmenting plaintext data and when segmenting encrypted sample data, respectively. The rather strange shape of the error concealment curve can be also attributed to the increasing number of complete segmentation failures. In any case, the low extent of overlap explains the very poor recognition results.

4. CONCLUSION

Although the amount of data to be compressed and encrypted when normalised texture is processed is much lower, it is by far more effective to partially encrypt iris sample data. It turns out that segmentation severely fails on encrypted sample data leading to proper protection when encrypting only 1% of the JPEG2000 packet data. Results also indicate that these findings do hardly depend on the employed recognition scheme, while for normalised texture encryption the resulting protection strength highly depends on the actual recognition scheme applied to the encrypted data. These results are in accordance with earlier findings in that segmentation is rather sensitive to image distortions like compression artefacts [15] and strongly determines recognition performance [16].

Based on these results, we have to reconsider earlier recommendations to apply selective JPEG2000 encryption to normalised textures [1]. The impact of applying different types of segmentation schemes is subject to future work.

5. REFERENCES

[1] M. Rieger, J. Hämmerle-Uhl, and A. Uhl, "Efficient iris sample data protection using selective jpeg2000 encryption of normalised texture," in *Proceedings of the 6th International Workshop on Biometrics and Forensics (IWBF'18)*, Sassari, Italy, 2018, pp. 1–7.

[2] Christian Rathgeb, Andreas Uhl, and Peter Wild, *Iris Recognition: From Segmentation to Template Security*, vol. 59 of *Advances in Information Security*, Springer Verlag, 2013.

[3] Christian Rathgeb, Andreas Uhl, Peter Wild, and Heinz Hofbauer, "Design decisions for an iris recognition sdk," in *Handbook of Iris Recognition*, Kevin Bowyer and Mark J. Burge, Eds., Advances in Computer Vision and Pattern Recognition. Springer, second edition edition, 2016.

[4] Dominik Engel, Thomas Stütz, and Andreas Uhl, "A survey on JPEG2000 encryption," *Multimedia Systems*, vol. 15, no. 4, pp. 243–270, 2009.

[5] F. Dufaux, S. Wee, J. Apostolopoulos, and T. Ebrahimi, "JPSEC for secure imaging in JPEG2000," in *Applications of Digital Image Processing XXVII*, Andrew G. Tescher, Ed. Aug. 2004, vol. 5558, pp. 319–330, SPIE.

[6] Martin Draschl, Jutta Hämmerle-Uhl, and Andreas Uhl, "Efficient fingerprint image protection principles using selective JPEG2000 encryption," in *Proceedings of the 1st Workshop on Sensing, Processing and Learning for Intelligent Machines (SPLINE 2016)*, Aalborg, Denmark, 2016, pp. 1–6.

[7] Martin Draschl, Jutta Hämmerle-Uhl, and Andreas Uhl, "Assessment of Efficient Fingerprint Image Protection Principles using different Types of AFIS," in *Proceedings of the 18th International Conference on Information and Communications Security (ICICS'16)*, Singapore, 2016, vol. 9977 of *Springer LNCS*, pp. 241–253.

[8] Martin Draschl, Jutta Hämmerle-Uhl, and Andreas Uhl, "Sensor dependency in efficient fingerprint image protection using selective jpeg2000 encryption," in *Proceedings of the 5th International Workshop on Biometrics and Forensics (IWBF'17)*, Coventry, United Kindom, 2017, pp. 1–6.

[9] Thomas Stütz and Andreas Uhl, "On JPEG2000 error concealment attacks," in *Advantages in Image and Video Technology: Proceedings of the 3rd Pacific-Rim Symposium on Image and Video Technology, PSIVT '09*, Tokyo, Japan, Jan. 2009, Lecture Notes in Computer Science, pp. 851–861, Springer.

[10] L. Ma, T. Tan, Y. Wang, and D. Zhang, "Efficient iris recognition by characterizing key local variations," *IEEE Transactions on Image Processing*, vol. 13, pp. 739–750, 2004.

[11] Libor Masek, "Recognition of human iris patterns for biometric identification," M.S. thesis, University of Western Australia, 2003.

[12] J.-G. Ko, Y.-H. Gil, J.-H. Yoo, and K.-I. Chung, "A novel and efficient feature extraction method for iris recognition," *ETRI Journal*, vol. 29, no. 3, pp. 399 – 401, 2007.

[13] D. Monro, S. Rakshit, and D. Zhang, "DCT-based iris recognition," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 29, no. 4, pp. 586–595, 2007.

[14] Dario Maio, Davide Maltoni, Raffaele Cappelli, James L. Wayman, and Anil K. Jain, "FVC2004: Third Fingerprint Verification Competition," in *ICBA*. 2004, vol. 3072 of *LNCS*, pp. 1–7, Springer Verlag.

[15] C. Rathgeb, A. Uhl, and P. Wild, "Effects of severe image compression on iris segmentation performance (best poster award)," in *Proceedings of the IAPR/IEEE International Joint Conference on Biometrics (IJCB'14)*, 2014.

[16] Heinz Hofbauer, Fernando Alonso-Fernandez, Josef Bigun, and Andreas Uhl, "Experimental analysis regarding the influence of iris segmentation on the recognition rate," *IET Biometrics*, vol. 5, no. 3, pp. 200 – 211, 2016.