Embedded Watermarking For Wireless Image Content Authentication

DOMINIC OSBORNE, DEREK ROGERS and DEREK ABBOTT
School of Electrical and Electronic Engineering and Centre for Biomedical Engineering (CBME)
The University of Adelaide
Adelaide, SA, 5005
AUSTRALIA
dosborne@eleceng.adelaide.edu

Abstract: - Wireless communication technology is one of the most rapidly expanding areas today, which has provided increased opportunity for end applications, such as teleradiology. This technology involves the electronic communication of diagnostic imagery for specialist interpretation, treatment and remote care. The cost of transmission bandwidth has been falling recently, however it is most cost-effective to compress wireless radiographic images for rapid and effective communication in these environments. Forward error correction techniques ensure that wireless data is transmitted robustly, but do not guarantee data integrity. A novel authentication watermarking technique is presented in this paper that extracts critical feature information from crucially important regions and embeds multiple robust watermarks into the image background. This provides a method of authenticating critical diagnostic feature information, while surviving acceptable levels of JPEG compression and the wireless communication channel.

Key-Words: - Authentication, Watermarking, JPEG, Compression, Robustness, Quality Assurance.

1 Introduction

The wireless transmission of diagnostic medical imagery for use in teleradiology has recently proven to be an efficient and cost-effective method for diagnosis and treatment to remote locations that would not otherwise receive assistance. This technology is still in its infancy and there remains a demand for implementation of useful quality assurance measures [1],[2]. Medical images, in particular, demand accurately known location and density information pertaining to the most essential features of the image called the Regions of Interest (ROI). This is because of the need to avoid erroneous diagnosis that could occur because of effects introduced by compression and transmission. Although the cost of transmission bandwidth is decreasing and effective techniques have been developed to protect data, there remains a strong need for authentication of these image types. Embedded watermarking is used to verify the integrity of medical images prior to any diagnosis that is made after transmission over a wireless link. There are two key benefits in using watermarking for the purposes of image content authentication. Watermarks eliminate the requirement of storing separate data, such as appended signatures and are designed to experience the same manipulations as the cover work. The watermarking technique is designed to survive acceptable distortions, such as JPEG compression, but fails when essential feature information from the ROI is lost. Best possible bit-rate performance is ensured by optimizing the watermarking parameters given an arbitrarily chosen mobile link and is compared to earlier methods in this field [3],[4]. Image integrity is provided with assurance that essential feature information has been retained.

2 Problem Statement

There are two types of distortions that may be applied to wireless diagnostic imagery. These include those encountered over the link, which may include the introduction of noise or small burst errors. The other types of distortions are legitimate types that are necessary for effective transmission, such as JPEG compression. This type of distortion is acceptable provided that only small levels are used [5],[6]. This can also be a problem if excessive levels are used that may result in a loss of diagnostic detail.

2.1 The Problem of Image Authenticity

Adding acceptable levels of noise to corrupt the bitstream of a JPEG image file that has been protected for the channel using Forward Error Correction (FEC) techniques does not usually result in loss of feature information in essential parts of the image. If many errors are encountered, there is usually a loss of image or obvious perceptual degradation. Incidental distortions that might not be corrected in the channel decoding process may slightly distort the file structure without any noticeable change to perceptual quality [7]. These methods do not guarantee image integrity. For example, the loss of a few image transform coefficients may incur a loss of detail that is not easily noticed. This is acceptable provided that the distortions introduced do not occur in the ROI. A particularly challenging example used is a Breast Tumor, in Fig. 1 that depicts this type of critical detail.

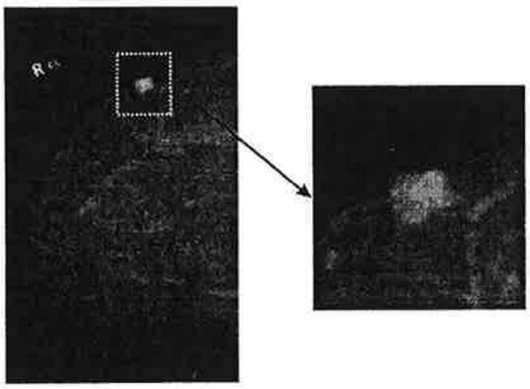

Fig 1: Encased carcinoma of the breast. This example highlights the importance of a known critical region containing most of the diagnostic content. Image degradation occurring in the region shown would be considered unacceptable as this is where all the diagnosis is being taken.

2.2 Minimizing the Bit Rate

To maintain diagnostic integrity using conventional methods, digital medical images are stored without loss of information using lossless compression schemes that do not discard image information. Mobile transmission of these image types for wireless telemedicine is prohibitive without using lossy compression techniques to reduce the image file sizes. For example, a typical mammogram may be digitized at 2048×2048 pixels at a bit-rate of 16 Bits Per Pixel (Bpp). This would result in a file exceeding 8 MB in size if compression is not used. On a 9.6 kilobit/s link this would require a transmission time of nearly two hours. The JPEG standard is the most widespread image compression technique, which is integrated into most digital industry-wide and consumer imaging equipment. Smart implementation of this standard is cost-effective as it will not compromise interoperability with other equipment. Recent research has shown two very interesting results. It has been demonstrated that acceptable levels of JPEG compression may be applied to the complete images without loss of quality [5],[6]. An acceptable threshold for use in medical imagery has been shown to be around 1.0 Bpp. Subjective testing in literature has shown that most medical images can be segmented into a ROI and a Region of Backgrounds (ROB), where the ROI contains all of the diagnostic information. For example, ROI segmentation and JPEG2000 compression was applied to 42 kidney nephrostograms to a level of 0.03 Bpp in [10]. Such methods permit greater levels of distortion in the ROB than for the ROI, making this region useful for the placement of authentication information. This ROI does not appear to exceed 20% of the area of most medical images shown in literature [8],[9],[10],[11].

3 Methods Used

3.1 Overview of Operation

If the authentication information is lumped and localized within the ROB and contains essential image information, it is possible to authenticate and verify the diagnostic integrity of such images. A useful approach involves extracting essential feature information from the ROI and embedding it in the ROB as depicted in Fig. 2

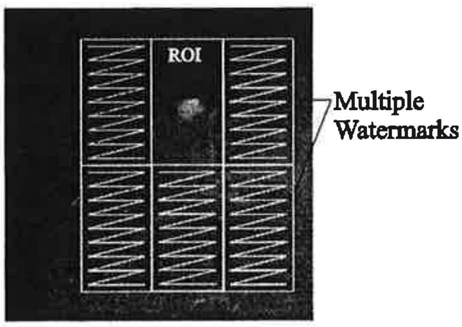

Fig. 2: The algorithm extracts a signature from the ROI and embeds it multiply into the micro blocks of the ROB. The number of watermarks used is subject to the availability of space and can be up to a maximum of 8.

Once the image is received, a signature can be extracted from the ROI and compared to the extracted watermark in the ROB. If the image had been clipped or degraded in an unimportant region, another watermark may be used for authentication. This assumes the receiver knows the location of the ROI. If excessive compression or losses have occurred in this region, the signature will not match any of the watermarks and will fail a test for authenticity. Excessive compression will also result in loss of watermark prior to a degradation of critical image feature information when JPEG quantization processes are applied. This is very essential as this information cannot be lost without incurring loss of a watermark.

3.1 Basis of Signature Extraction

The basis of singular watermark extraction and embedding has been adopted from a method proposed by Lin and Chang [4]. This method has been designed to withstand variable levels of JPEG compression and is dependent on the watermark embedding strength used. The signature extracted is based on the non-varying properties of the Discrete Cosine Transform (DCT) coefficients that do not change after an acceptable level of JPEG quantization has been applied. JPEG compression involves performing the DCT on 8×8 image pixel blocks to create micro blocks. Two steps in the JPEG compression process reduce the number of bits representing an image. These include quantization and rounding of the DCT coefficients and entropy coding. The first two steps are a lossy operation, which change pixel values but maintain important visual properties of the image. The third stage does not result in loss to the image. Let $x \bullet q$ be the result of quantizing x to an integral multiple of a quantization step size, q:

$$x \bullet q = q\left\lfloor \frac{x}{q} + 0.5 \right\rfloor . \tag{1}$$

Consider s to be a real valued scalar quantity used to represent a DCT coefficient and q_1 and q_2 as quantization step sizes with $q_2 \leq q_1$, then:

$$((s \bullet q_1) \bullet q_2) \bullet q_1 = s \bullet q_1 . \tag{2}$$

Hence if s is quantized by the larger step, q_1 and then by a smaller step, q_2, the second quantization step can be reversed. The watermark should survive provided that the quantization performed during compression uses smaller step sizes. Because all of the micro blocks are divided by the same quantization table in the JPEG compression process, the relationship between two DCT coefficients located at the same coordinate position does not change after quantization. This property allows features from the most important DCT coefficients to be extracted from the ROI to be used as a signature that can be embedded into the ROB. Micro blocks from the ROI are grouped into pseudo-random pairs using a seed. For each pair of blocks, 8 corresponding DCT coefficients are compared to obtain 8 bits of a binary signature. These represent the most significant information in the ROI as the remaining coefficients are typically zero after JPEG quantization. Mid to upper frequency DCT coefficients are used in the ROB for the embedding of the signature as a watermark. Consider two blocks that have been grouped, C_a and C_b, where i and j are the coordinates of a signature coefficient, K_i then:

$$K_i = \begin{cases} 0 \ : \ C_a[i,j] < C_b[i,j] \\ 1 \ : \ C_a[i,j] \geq C_b[i,j]. \end{cases} \tag{3}$$

The coefficients used for the signature and watermark are illustrated in Fig. 3

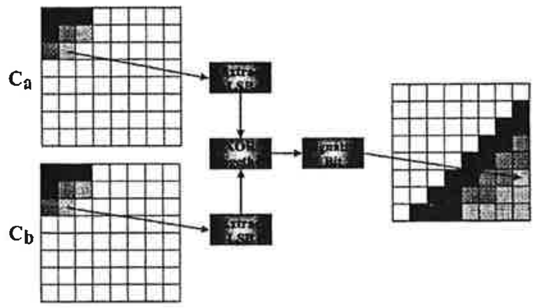

Fig 3: The left blocks of this diagram depicts the DCT coefficients used to extract a signature bit involving the exclusive-or of the Least Significant Bit (LSB.) This is one way of extracting a signature from these coefficients. Matching shades indicate where the payload bits go in the ROB. This approach is different to that proposed by Lin and Chang [4], who embedded authentication information in the same micro blocks where signature information was extracted.

3.2 Watermark Embedding

The method used to embed the signature requires an additional embedding strength parameter, α to be specified and is similar to an implementation of the method proposed by Lin and Chang [4]. This specifies the level of JPEG compression that the watermark should survive and can be expressed as the largest multiplier that can be applied to the quantization matrix which results in the positive detection of a watermark. Four signature bits are embedded into the each micro block in the ROB although this could be increased if a larger watermark capacity is required. For each signature bit ρ, seven coefficients are selected from the right side of Fig. 3. The JPEG compression operation is used on this coefficient group where each coefficient is quantized by its corresponding quantization factor α, and rounded to the nearest integer as of Equation 4, where $Q[i]$ is the corresponding quantization value.

$$C_I[i] = \mathrm{ROUND}\left[\frac{C[i]}{\alpha Q[i]}\right]. \quad (4)$$

The maximum level of acceptable distortion is set by the appropriate choice of α, which is also the embedding strength used. The LSB of the integers C_I are exclusive-OR'd together to obtain the present bit value, ρ_e. If $\rho_e \neq \rho$, then the LSB of one of the integers is flipped. The resulting watermarked coefficients are multiplied by their corresponding quantization factors to generate the watermarked DCT coefficients:

$$C_{DCT}[i] = \alpha Q[i] C_w[i]. \quad (5)$$

Acceptable distortions occur if the image is converted to its spatial format, or if different decimation algorithms are used in the JPEG compressor. Clipping and rounding 8-bit values has the effect of introducing noise into the watermarking system. Experiments for this paper found that this did not exceed 15%. It is thus very useful to use a tolerance bound, τ to the number of correct watermark bits extracted from the ROB to those found in the original signature in the ROI. This is desired when these bits are compared pair-wise in the authentication process. For the experiments conducted, τ = 15%. Hence the percentage of matching bits must exceed 85% for an authenticated image. Different watermarks can be used for this match, the best match being chosen. At least one watermark must be present for authentication.

3.3 Image Restoration

The ROI watermark may be removed from the authenticated image prior to diagnosis. The watermarked region in the ROB can be approximately restored by setting the high frequency DCT coefficients to zero, as shown on the right side of Fig. 4. This occurs in each micro block to eliminate watermarking artifacts.

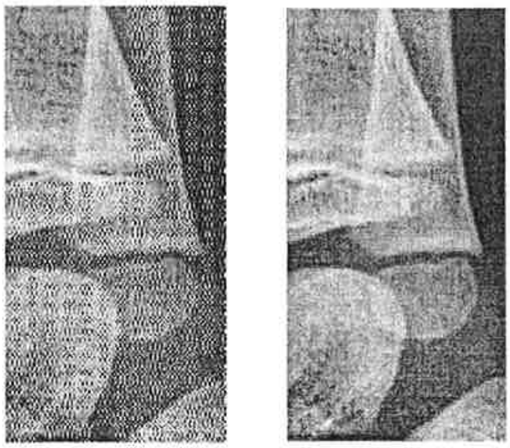

Fig 4: Removal of watermark that is present in all regions in the image on the left using an embedding strength, $\alpha = 2.5$. This level is unusually large but is used to illustrate the watermark artifacts. An exception is the top left corner that contains no watermark as it is part of the ROI. The image on the right has had the watermark removed. Close inspection of these artifacts shows that they take the form of checkerboard patterns. This is a typical result of embedding into the DCT image domain.

4 Comparison of Methods

The performance of the ROI watermarking scheme is compared to an earlier method proposed by [4]. This approach is designed to operate on complete images rather than targeting specified regions. The tests that are used for a comparison include watermark robustness to JPEG compression with the application of DCT quantization as well as complete baseline JPEG and lossy JPEG-2000 compression. The Peak Signal-to-Noise Ratio (PSNR) is used to compare the performance of the image fidelity with the application of watermarking in both systems over a similar range of embedding strengths. In order to choose appropriate embedding strengths a sample set of 500 grayscale medical images with resolutions between 512×512 and 2048×2048 pixels are used to determine an average relationship between the applied quantization level and the resulting compression level. This makes it possible to select appropriate levels of quantization for the desired bit rate of 1.0 Bpp, shown in Fig. 5. The ROI is consistently specified to occupy 20% of the medical image and is located at the image center, for ease of calculation so that 8 surrounding watermarks can be applied.

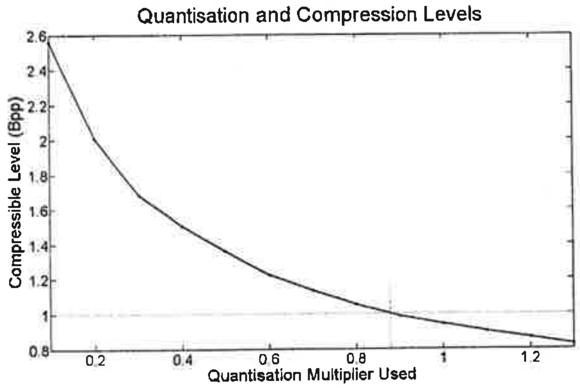

Fig. 5: Relationship between compression and quantisation level. In order to achieve the desired bit rate of 1.0 Bpp, a minimum embedding strength, $\alpha = 0.88$ is required.

4.1 Robustness

The percentage of images detected as authentic are used as an indication of the degree of robustness of the watermarking schemes to JPEG compression. A comparison between the ROI watermark method and that proposed by Lin and Chang is shown in Fig. 6. The system is specified to fail when the percentage of authentic images is less than 85%.

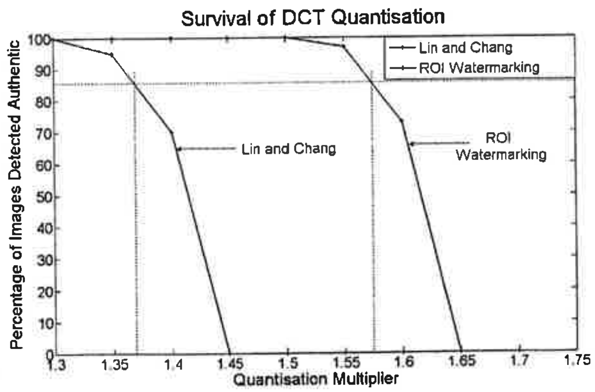

Fig. 6: Comparing the survival of the watermark subject to applied DCT quantization for the ROI watermark method and the method proposed by Lin and Chang [4], using 500 grayscale medical images.

For each of these watermarking systems robustness to DCT quantization far exceeds the embedding strength, $\alpha = 0.88$. Lin and Chang's method fails when the quantization multiplier exceeds 1.37, while the ROI watermark approach fails when this level exceeds 1.58. Thus the quantization level can be exceeded by 0.21. This would suggest that increased levels of JPEG compression could be applied, while maintaining watermark integrity and hence important features in the image.

4.2 Compression Performance

Two compression schemes are used to evaluate the effectiveness of the ROI watermarking scheme to survive the target bit rate of 1.0 Bpp. These include baseline JPEG and JPEG-2000 compression standards. Previous testing found that the method survived acceptable levels of DCT quantization. This was a specific baseline JPEG distortion, but did not include the entire set of distortions encountered during the complete JPEG compression process. Survival to JPEG-2000 compression scheme is treated as a side issue as it is wavelet-based, as watermarking takes place in the DCT image transform domain. JPEG-2000 compression testing is used to investigate the potential for both compression systems to be used at the low target bit rate, yet maintain survival of the watermark. Images were watermarked with $\alpha = 0.88$ corresponding to a level expected to survive baseline JPEG compression of 1.0 Bpp. For the purposes of conventional DCT-based JPEG compression, a compressor is used from the Matlab Image Processing Toolbox. Precision is limited to 0.05 Bpp for the baseline JPEG compressor. The type of JPEG-2000 compression used is an implementation of Part 1 of the JPEG-2000 standard. This software was developed by David Taubman [13] and is an implementation called Kakadu. Irreversible lossy JPEG-2000 compression is used and the bit rate specified to the compressor directly as a floating point number. Results are shown in Fig. 7.

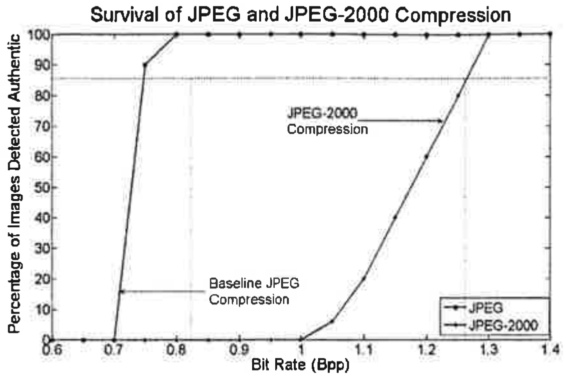

Fig. 7: Comparing the survival of the ROI watermark to baseline JPEG and lossy JPEG-2000 compression schemes. It is apparent that survival to JPEG-2000 compression fails when the compression level is between 1.2 and 1.3 Bpp. Hence the watermarking method fails to authenticate JPEG-2000 images compressed to target bit rates of 1.0 Bpp. To the contrary, survival to baseline JPEG compression is a success as the watermark can be compressed 0.3 Bpp past the expected survival rate of 1.0 Bpp.

4.3 Image Fidelity

The PSNR is used as an indication of the degree of watermarking artifacts present in the overall image. A larger value represents more similarity to an unwatermarked work, which is favored. The embedding strength is specified as the greatest multiplier that can be applied to the quantization matrix without loss of a watermark. A comparison is shown in Fig. 8:

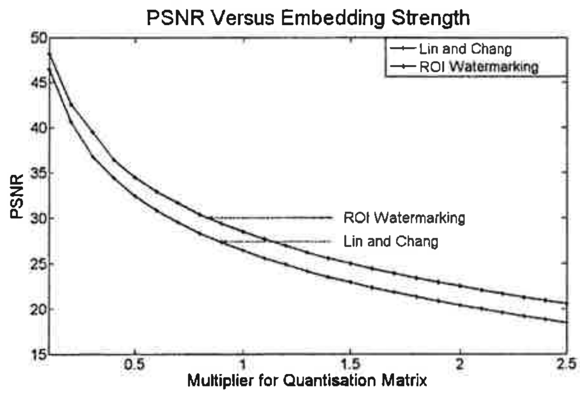

Fig. 8: ROI watermarking appears to show marginal improvement over Lin and Chang's method. The most likely reasons are less area being devoted to the watermark, as it does not need to occupy the entire image region. PSNR is shown to be strongly dependent on the embedding strength, α. Increasing the embedding level results in a lowered PSNR. This is a result of increased watermarking artifacts that elevates the difference between watermarked and unwatermarked images as α is increased.

5 Discussion

The results of testing the performance of the ROI watermarking system are promising. They show that the technique outperforms earlier methods with improved robustness to compression and image fidelity. Resilience to JPEG compression shows that the watermark will remain embedded in the image after increased levels of JPEG quantization are used if the ROI watermark system is adopted rather than earlier methods. This would appear to be a significant result, as greater than anticipated signal distortions may be applied without loss to essential feature information contained in the ROI. Improved quantization performance would appear to be a result of adopting a ROI watermark approach, rather than embedding into all image regions. A smaller signature would suggest that fewer bits can be corrupted and this is reflected in increased robustness. Watermarking is also much more efficient as fewer bits are required to authenticate the important region of a medical image. Image fidelity has shown to be slightly improved over earlier methods, with larger PSNR at all embedding levels used. Some of the experiments could have been tested in ways that may have given more realistic or slightly different results. Both the quantization and compression tests have specified the ROI to be rectangular and based at the center of the test image occupying 20% of the image area. This had the benefit of allowing 8 watermarks to be embedded around the ROI, but did not provide for circumstances when this region was located near the image edges. This could have reduced the number of watermarks in test situations where the ROI was placed in an image corner, which would have resulted in 3 watermarks. One of the pressing questions is whether having a fewer average number of watermarks would be expected to change the results significantly? Both the robustness and compression tests have demonstrated improvements in performance over earlier methods, however this has most likely been a result of a smaller authentication signature leading to fewer bit errors between the ROI features and the watermark after distortions had been applied. As all of the watermarks are expected to degrade at a similar quantization or compression level, having a different number of watermarks by changing the position of the ROI in the image may not significantly change the existing results when the ROI is set to the image center. One way of rectifying this experimental limit would be to set the position of the ROI as a random variable containing of all of the possible values that one location within the ROI could take. This position could be used as a reference point. Ideally the size of the ROI should have been specified to vary within some acceptable range or limit. One last area that could have been improved would be the method to compare the signature and watermark. These sequences were compared on a bitwise pair basis. A preferred option may have been to determine the cross correlation between the signature and watermark. This would have the capability to compensate for errors involving a loss of part of the signature or watermark. Fortunately this type of loss is highly improbable, as errors resulting from minor image distortions typically only cause bits to flip their value.

6 Future Work

To verify the integrity of the ROI DCT coefficients in the most complete way possible would involve extracting all of the bits that represent significant picture information for the payload to a surrounding watermark. Assuming that 10 bits are used to represent each DCT coefficient (including a sign bit) and 8 coefficients are used to extract a signature, this would require 80 bits to be extracted from every micro block. It would be particularly challenging to attempt to embed this extent of information given a limited watermark capacity. If it were desired to increase the watermark capacity so that greater amounts of signature information could be placed into each micro block of the ROB, it may be possible to modify the algorithm in two ways. Firstly the number of coefficients used to contain an embedded bit could be reduced from 7 to 1 coefficient per bit. This would increase the watermark capacity from 4 bits per micro block to 28 bits per block. Secondly, if a greater watermark capacity was still required, it may be possible to embed more than one bit into a coefficient. This could be made possible by flipping more than one bit in the binary representation of a coefficient. The benefit of significantly greater watermark capacity would probably be outweighed by reduced survival and robustness to compression-related distortions.

Another area of this work that may require further scrutiny is the approximation used to select quantization table multipliers for this system. In Fig. 5, a relationship was established between the amount of quantization used and the average level of expected compression obtained. This provides some approximations to appropriate levels of quantization and embedding strength parameters that can be used to optimize system performance over the target bit rate of 1.0 Bpp. The alternative and more accurate approach may be to analyze individual images to determine the compression level resulting from variable quantization levels. This is because the compression-quantization relationship is specific to the image used. A broad selection of image types and resolutions are chosen so that an average quantization-compression curve is determined for an approximation to the embedding strength. Determining this parameter for individual images would be expected to increase the overall system complexity by increasing the number of computations required, but may be preferred for a more accurate choice of α. This factor is shown to be dependent on the watermark failure point at a given compression level, where greater α leads to increased watermark survival. This problem of ensuring watermark failure after this level has been exceeded may have been underestimated as it is equally important that the watermark fail when unacceptable or excessive compression levels are used. It may be more appropriate to optimize the algorithm for individual images. This could be expected to give a more accurate choice of α, which would allow for the point of watermark failure to be predicted more easily. The only trade-off to make would be a slight increase in system complexity and time to encode the watermark resulting from a repetition of watermark embedding processes that would otherwise only be implemented once.

7 Conclusion

Wireless telemedicine has recently proven to be an efficient and cost-effective method for diagnosis and improved specialist medical treatment to remote locations. This technology demands effective standards and quality assurance measures to be available for implementation to the user. This is especially important for the end application of teleradiology that involves the transmission of wireless medical images. Such standards are considered as a sign that the technology has reached full maturity and can be designed, implemented and deployed with less concerns for the integrity of transmitted data. Research focus has been to ensure image quality detailing embedded watermarking to verify the integrity of critically important diagnostic features, given the challenges and constraints of a wireless communications system. Careful attention has been given in the design to minimize costs by facilitating the re-use of existing communications infrastructures and widespread consumer and specialist image compression standards. Cellular technology has been shown to be one of the most practical and feasible options when used in conjunction with baseline JPEG image compression, as the research focus is based on minimizing the bit rate. Wireless transmission of images used in teleradiology is shown to be most rapid and effective when maximum acceptable compression levels are used. The authentication watermarking system presented is shown to maintain compatibility with JPEG. This work has optimized the speed of transmission by ensuring that the medical image authentication system is suited to the best possible bit rate, given the maximum permitted image compression level. Attention has been given to minimize degradation in diagnostically significant regions by embedding authentication information for two-level verification into the ROB. Performance testing of the authentication watermarking system has been promising. The method presented outperforms earlier techniques with improved robustness, compression performance and image fidelity. Resilience to JPEG compression shows that the watermark will remain embedded in the image after increased levels of JPEG quantization are used if the ROI watermark system is adopted rather than earlier methods. Improved robustness would appear to be a result of adopting a ROI watermark approach, rather than embedding into all image regions. A smaller signature would suggest that fewer bits can be corrupted and this is reflected in an increased robustness. Image fidelity has shown to be significantly improved over earlier methods, with improved PSNR at all watermark embedding strengths used.

References:

[1] D. Osborne, D. Abbott, M. Sorell. and D. Rogers. Embedded Importance Watermarking For Image Verification In Radiology, *SPIE International Symposium - Microelectronics, MEMS and Nanotechnology*, Vol.4935, 2003, pp. 383-390.

[2] R. D. Silverman. Current Legal and Ethical Concerns in Telemedicine and E-Medicine, *Journal of Telemedicine and Telecare*, Supplement 1, 2003, pp 67-69.

[3] G. K. Wallace. The JPEG Still Image Compression Standard. *Communications of the ACM* 34,4, April 1991, pp. 30-45.

[4] C. Y. Lin. and S. F. Chang. A Robust Image Authentication Method Distinguishing JPEG Compression from Malicious Manipulations, *IEEE Transactions on Circuits and Systems for Video Technology*, Vol 11, No 2, 1998, pp. 153- 168.

[5] R. M. Slone, E. Muka and T. K. Pilgram. Irreversible JPEG Compression of Digital Chest Radiographs for Primary Interpretation: Assessment of Visually Lossless Threshold, *Radiology*, Vol 228, 2003, pp. 425-429.

[6] O. Kocsis, L. Costaridou, G. Mandellos, D. Lymberopolous and G. Panayiotakis. Compression Assessment Based on Medical Image Quality Concepts Using Computer Generated Test Images, *Computer Methods and Programs in Biomedicine*, Vol 71, No 2, 2003, pp. 105-115.

[7] A. Viterbi. Error Bounds For Convolutional Codes and an Asymptotically Optimum Decoding Algorithm. *IEEE Transactions on Information Theory* IT-13, April 1967, pp. 260-269.

[8] M. B. Kim, Y. D. Cho, D. K. Kim. and N. K. Ha. On the Compression of Medical Images With Regions of Interest. *SPIE Visual Communications and Image Processing*, Vol 2501, September, 1995 pp. 733-744.

[9] M. Penedo, W. A. Peraman, P. G. Tahoces, M. Souto and J. J. Vidal. Embedded Wavelet Region-Based Coding Methods Applied to Digital Mammography, *IEEE International Conference on Image Processing*, September, Vol 2, 2003, pp. 197-200.

[10] G. K. Anastassopoulos and A. Skodras. JPEG2000 ROI Coding in Medical Imaging Applications, *Proceedings of the 2nd IASTED International Conference on Visualisation, Imaging and Image Processing*, September, 2002, pp. 783-788.

[11] A. Wakatani. Digital Watermarking for ROI Medical Images by Using Compressed Signature Image, *Proceedings of the 35th Hawaii International Conference on System Sciences, IEEE Computer Society*, Vol 6, 2002, pp. 157- 163.

[12] I. J. Cox, M. L. Miller and J. A. Bloom. *Digital Watermarking*, Morgan Kaufmann Publishers, 2002.

[13] D. S. Taubman and M. W. Marcellin. JPEG-2000: Image Compression Fundamentals, Standards and Practice, Kluwer Academic Publishers, 2001.