Watermarking digital images for copyright protection

J.J.K.Ó Ruanaidh
W.J.Dowling
F.M.Boland

Indexing terms: Copyright protection, Image processing, Steganography, Spread spectrum communications

Abstract: A watermark is an invisible mark placed on an image that is designed to identify both the source of an image as well as its intended recipient. The authors present an overview of watermarking techniques and demonstrate a solution to one of the key problems in image watermarking, namely how to hide robust invisible labels inside grey scale or colour digital images.

1 Introduction

Computers, printers and high rate transmission facilities are becoming less expensive and more generally available. It is now feasible and very economical to transmit images and video sequences using computer networks rather than to send hard copies by post. In addition, images may be stored in databases in digital form. A major impediment to the use of electronic distribution and storage is the ease of intercepting, copying and redistributing electronic images and documents in their *exact* original form. As a result, publishers are extremely reluctant to use this means of disseminating material. The commercial possibilities for the World Wide Web are steadily becoming more appreciated. However, if these possibilities are to be realised, an integrated approach to the secure handling, issue and duplication of issued documents is required. Public key encryption systems such as the RSA algorithm [1–3] do not completely solve the problem of unauthorised copying because of the ease with which images may be reproduced from previously published documents. All encrypted documents and images need to be decrypted before they can be inspected or used. Once encryption is removed the document can be passed on in an electronic form. If there is more than one recipient of an image, there is no direct proof that any particular authorised recipient is responsible for passing it on to unauthorised users. The idea of using an indelible watermark to identify uniquely both the source of an image and an intended recipient has therefore stimulated much interest in the electronic publishing and printing industries.

To be effective, an embedded watermark should be visually imperceptible, secure, reliable and resistant to attack.

Imperceptible. The image must not be visibly degraded by the presence of the mark. The mark should serve as a unique identifier with a high information content.

Secure and reliable. The mark must be strongly resistant to unauthorised detection and decoding. The watermark must also be capable of identifying the source and intended recipient with a low probability of error. It is also desirable that it would be difficult for an unauthorised agent to forge watermarks. Innovative error-control coding and digital signature techniques are required to ensure reliable and secure communication of the mark as well as authentication of the encoded message.

Robust. The mark must be robust to attack and must be tolerant to reasonable quality lossy compression of the image using transform coding, vector quantisation or any other technique. Standard image processing operations such as low pass filtering, cropping, translation and rescaling should not remove the mark.

Later we shall describe a method which fulfils most of the above requirements. In this paper, we argue that watermarking needs to be *adaptive* in order to be robust. In direct contrast to many other techniques, with the notable exception of Cox *et al.* [4], the method here places the watermark on the *most perceptually significant* components of an image. The logic behind the premise is quite simple. A watermark that is nonintrusive is one which resembles the image it is designed to protect. By virtue of its similarity to the image, any operation that is intentionally performed to damage the watermark will also damage the image.

The factors affecting the transmission of information embedded in images are quite complex. First, there is the need for robustness. The second factor is visibility. Intuitively, one can see that less information can be hidden on flat featureless regions of the image. It should be possible to incorporate more information into those parts of the image that contain more texture or around edges, provided edge integrity is maintained. Psychovisual phenomena are obviously factors in the transmission of hidden information.

There are two main principles involved in designing a watermark. The first principle, mentioned earlier, is

IEE Proceedings online no. 19960711

Paper first received 22nd December 1995 and in revised form 14th June 1996

J.J.K. Ó Ruanaidh was with Trinity College Dublin and is now with the Computer Vision Group, Centre Universitaire d'Informatique, 24 Rue Général Dufour, Université de Genève, CH 1211 Genève 4, Switzerland

W.J. Dowling and F.M. Boland are with the Department of Electronic and Electrical Engineering, Trinity College Dublin, Dublin 2, Ireland

that a successful watermarking algorithm should explicitly identify and place the mark in the most important features of the image. There are some similarities to the key ideas behind image compression and there will be many ideas and techniques borrowed from this field. The second principle, which we shall outline briefly, is that of *spread spectrum communications* [5].

2 Previous work

Brassil *et al.* [6] have investigated different methods for marking text within documents with a unique binary codeword which serves to identify legitimate users of the document. The codeword is embedded in a document by making subtle modifications to the structure of the document such as modulation of line width and interword spacing as well as modification of character fonts. The presence of the codeword does not visibly degrade the document but can be readily detected by making a comparison with the original. Standard document handing operations such as photocopying and scanning do not remove the mark. The same idea may be extended to include the protection of images.

Kurak and McHugh [7] have considered the possible application of redundant features in digital images to the transmission of information. Their concern was the transmission of dangerous viruses (or 'Trojan horse programs') in the least significant bits of a data stream. They note that merely viewing an image is not sufficient for detecting the presence of some form of corruption. Depending on the texture of the image and the quality of a computer monitor, it is possible to exploit the limited dynamic range of the human eye to hide low-quality images within other images. Walton [8] has developed a technique for introducing checksums in the least significant bits of an image to implement a fragile watermark and thus prevent unauthorised tampering. Dautzenberg and Boland [9] examined the use of the least significant bits as a possible scheme for introducing watermarks into images. This approach gave very poor results because standard lossy compression schemes, such as JPEG [10], tend to have the effect of randomising the least significant bits during the quantisation stage of image compression.

Zhao and Koch [11] have investigated an approach to watermarking images based on the JPEG [10] image compression algorithm. Their approach is to segment the image into individual 8×8 blocks. Only eight coefficients occupying particular positions in the 8×8 block of DCT coefficients can be marked. These comprise the low frequency components of the image block, but exclude the mean value coefficient (at coordinate (0,0)) as well as the low frequencies at coordinates (0,1) and (1,0). Three of the remaining DCT coefficients are selected using a pseudorandom number generator to convey information. The resemblance of this technique to frequency hop spread spectrum communications is mentioned by the authors [11]. Zhao and Koch also take the precaution of placing the blocks at random positions in the image in order to make a successful attack by an enemy less likely.

Tirkel *et al.* [12, 13] and van Schyndel *et al.* [14, 15] have applied the properties of m-sequences to produce watermarks that are resistant to filtering, image cropping and are reasonably robust to cryptographic attack. The original image is not required to decode the mark. Recent work [15] indicates progress towards producing more robust watermarks.

Matsui and Tanaka [16] have applied linear predictive coding for watermarking video, facsimile, dithered binary pictures and colour and grey scale images. Their approach to hiding a watermark is to make the watermark resemble quantisation noise. To a certain extent, their approach can be considered to be perceptually adaptive because quantisation noise is concentrated around edges and textured features. Cox *et al.* [4] believe that this method may not be robust to cropping. Ó Ruanaidh *et al.* [17] and Cox *et al.* [4] have developed perceptually adaptive transform domain methods for watermarking. In direct contrast to the previous approaches listed above the emphasis was on embedding the watermark in the *most significant* components of an image. The general approach used in these papers is to divide the image into blocks. Each block is mapped into the transform domain using either the discrete cosine transform [10, 18–20], the Hadamard transform [1, 18] or the Daubechies wavelet transform [19]. Only the components that are most significant to image intelligibility are marked. A transform-based watermarking algorithm is described in more detail in Section 4.

Transform domain modulation schemes possess a number of desirable features. First, one can mark according to the perceptual significance of different transform domain components which means that one can adaptively place watermarks where they are least noticeable, such as within the texture of an image. As a result, a transform domain watermark tends to resemble the original image. The watermark is also irregularly distributed over the entire image sub-block which makes it more difficult for enemies in possession of independent copies of the image to decode and to read the mark.

The scheme described by Cox *et al.* [4] differs from that used by Ó Ruanaidh *et al.* [17] in several ways. The main differences lie in the detection and decoding of the mark. Cox *et al.* embed a unique Gaussian distributed sequence into the coefficients. The Gaussian distribution is chosen to prevent attacks by colluding parties comparing independent copies of the image. Ó Ruanaidh *et al.* employ an alternative approach whereby a binary code is directly embedded in the image. One advantage of the latter approach is that it avoids the need to maintain large databases of watermarks. A disadvantage is that the sequences thus produced are discrete valued and therefore the watermark is less resistant to colluding parties. However, there is nothing to prevent one from using continuous watermarks to convey digital information. This would combine the best features of both approaches.

The discrete Fourier transform (DFT) may also be used in watermarking. The discrete Fourier transform of a real image is generally complex valued. This leads to a magnitude and phase representation for the image. Transform domain methods described above mark the components of real valued transforms. Ó Ruanaidh *et al.* [21] and Ó Ruanaidh *et al.* [17] have also investigated the use of DFT phase for the transmission of information. There are a number of reasons for doing this. First and most importantly, the human visual system is far more sensitive to phase distortions than to magnitude distortions [22]. Oppenheim and Lim [23] investigated the relative importance of the phase and magnitude components of the DFT to the intelligibility of an image and found that phase is more significant.

Second, from communications theory, it is well known that phase modulation can possess superior noise immunity when compared to amplitude modulation.

3 The block-mean approach

In this Section, we present an algorithm that forms the basis of understanding more sophisticated transform domain algorithms described later. Dautzenberg and Boland [9] and Caronni [24] have investigated a very simple technique for embedding watermarks in images. An image is divided up into blocks. The mean of each block may then be incremented to encode a '1' or decremented to encode a '0' (or vice versa). This is termed bi-directional coding. Alternatively, the mean may be incremented to encode a '1' and left unchanged to encode a '0'. This is termed unidirectional coding.

The block-mean approach suffers from the grave disadvantage that an enemy that is in possession of a number of independent copies of the image can compare the different copies and read most, if not all, of the encoded message. Caronni [24] shows that the expected number of undetected bits decreases exponentially with the number of copies. Caronni combats this particular weakness by randomising both the size of the blocks as well as the positions of the blocks inside the image. Despite its simplicity, the block-mean method of marking images has proven to be highly robust to lossy image compression, photocopying and colour scanning and dithering [24, 9].

The number of bits that may be encoded using the block-mean approach equals the number of blocks, and this in turn depends on the size of the image and the block size, as well as the width of borders around blocks. Realistically, for a typical image of size 256 × 256 pixels the number of bits that one can expect to encode is approximately one hundred bits. This number of bits may be adequate for some applications, even after taking into account the need for redundancy in the code for error detection and correction as well as code word authentication. However, as we will see, this capacity may be greatly increased by watermarking in the transform domain.

4 Transform domain watermark

This Section describes a transform domain watermarking algorithm. First, a simple form of modulation for placing bits on an image is outlined. Secondly, a technique for determining the number of bits to be placed at given locations in the image is described. Note that in this Section the DCT will be applied exclusively. However, there is no reason why other transforms cannot be applied in its place. Indeed, later in the paper examples of the use of other transforms for watermarking will be presented.

The following algorithm, which is adapted from JPEG [10] image compression and which is a hybrid between amplitude modulation and frequency shift keying, has been applied to watermarking:

1. Divide the image into blocks.

2. Subtract the mean of the block from each pixel in the block.

3. Normalise pixel values within each block so that they range between −127 and 127.

4. Compute the transform of the image block.

5. Modulate selected coefficients of the transformation (e.g. using bidirectional coding). The coefficients that are selected are those that are *most* relevant to the intelligibility of the image.

6. Compute the inverse transform, denormalise, add the mean to each pixel in the block and replace the image block in the image.

Steps 2 and 3 above produce a normalised image subblock with zero mean. Although one of the DCT coefficients computed in Step 4 already contains the mean, Step 2 is not redundant because the normalisation in Step 3 may only be carried out if the mean of the block is zero.

Watermark detection is easily performed by carrying out Steps 1 to 4 above on the original image and the watermarked image in parallel and comparing the values of the coefficients.

4.1 The number of bits

The most important factor in embedding a bit stream in an image is to determine the number of bits that can be placed into a given image block.

In a highly textured image block, energy tends to be more evenly distributed among the different DCT coefficients. In a flat featureless portion of the image the energy is concentrated in the low frequency components of the spectrum.

As stated earlier, the aim is to place more information bits where they are most robust to attack and are least noticeable. This may be accomplished by using a simple thresholding technique. The first stage is to use visual masking and to weight the transform coefficients $F(k_1, k_2)$, $0 \leq k_1 < N_1$ and $0 \leq k_2 < N_2$, according to a subjective measure of their visual perceptibility:

$$G(k_1, k_2) = w(k_1, k_2)F(k_1, k_2) \tag{1}$$

The most significant components are then selected by comparing the component magnitude squared to the total energy in the block. The coefficient $F(k_1, k_2)$ is selected if

$$|G(k_1, k_2)|^2 \geq \epsilon \sum_{k_1=0}^{N_1-1} \sum_{k_2=1}^{N_2-1} |G(k_1, k_2)|^2 \tag{2}$$

The quantisation tables [10] used in JPEG image compression can be exploited to choose the weighting in eqn. 1 for DCT watermarking with 8 × 8 blocks.

Lossy image compression algorithms are designed to disregard redundant information. Information bits placed within textured areas of the image are therefore more vulnerable to attack. There is a compromise to be reached between hiding a large number of information bits where they can least be seen, but where they can be attacked by image compression algorithms, or placing fewer bits on less textured but safer portions of the image. This may be achieved by opting for a moderately low value of threshold (e.g. $\varepsilon \approx 0.2$).

It is worth noting that the number of bits that can be encoded using image transforms far exceeds that of the block-mean approach. The number of modulated DCT coefficients is generally around 10 000 for a typical image. In the case of Zhao and Koch's method, 3 bits of information are encoded into each 8 × 8 block. If the blocks are tiled over the image then one could obtain a maximum code rate of 3/64 bits/pixel.

It is important to note the differences between the aims in image compression and in watermarking

images. In transform-based image compression, the goal is to obtain a small number of transform coefficients which can be used to obtain a good approximation to the original image. Small changes in the coefficient values should make little difference to the reconstructed image. However, the reverse does not necessarily hold since a small change to the image can result in a large change in the coefficient values (particularly when the basis images also change). This behaviour is obviously extremely undesirable since the embedded information depends on the value of these coefficients. The severity of this effect depends on the image transforms being used. Ill-conditioning tends to be much more severe for image transformations whose basis images are data-dependent (e.g. the singular value decomposition (SVD)). Image transformations with fixed basis functions (e.g. DCT and wavelet transforms) tend to exhibit more stable behaviour.

5 Reliable communications

The material in this paper thus far has described methods for watermarking images. However, we have not yet addressed the other main component in the watermarking problem, namely the reliable transmission of the watermark.

Reliable communication was proven by Shannon [25] to be theoretically possible providing the information rate does not exceed a threshold known as the channel capacity. In this Section we make some rather idealised assumptions regarding the form of the noise **n** corrupting a watermark and use information theory to derive rules for setting the optimal strength and location of the watermark **x**.

Let us write,

$$x_i + n_i = y_i \qquad 1 \leq i \leq N \tag{3}$$

where x_i is one element of a watermark vector of length N, n_i is an element of a *noise* vector and y_i is a element of a watermark distorted by image processing noise. All forms of image processing including vector quantisation, filtering and scanning introduce noise which degrades the watermark. We assume that the noise is additive, white, stationary and Gaussian:

$$p(y_i|x_i) = p(n_i) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left[-\frac{(y_i - x_i)^2}{2\sigma^2}\right] \tag{4}$$

We also assume that the n_i are uncorrelated and that

$$p(y_1, y_2 \cdots y_N|x_1, x_2 \cdots x_N) = \prod_{i=1}^{N} p(y_i|x_i) \tag{5}$$

Channel capacity [26] may be defined as

$$C = \max_{p(x)} I(\mathbf{X}; \mathbf{Y}) \tag{6}$$

where the watermark probability density function $p(x)$ is chosen to maximise the average mutual information $I(\mathbf{X}; \mathbf{Y})$.

According to Proakis [27] the capacity is maximised with respect to the distribution $p(x)$ if

$$p(x_i) = \frac{1}{\sqrt{2\pi\gamma^2}} \exp\left[-\frac{x_i^2}{2\gamma^2}\right] \tag{7}$$

which is a zero mean Gaussian density with variance γ^2. In this case,

$$I_{\max} = \frac{1}{2} N \log_2\left[1 + \frac{\gamma^2}{\sigma^2}\right] \tag{8}$$

Note that eqn. 7 would seem to support the use of a Gaussian distributed watermark such as that used by Cox *et al.* [4].

In image watermarking we might expect that the transmission of information is functioning under quite extreme conditions, in which case $\sigma^2 >> \gamma^2$, which implies

$$\ln\left(1 + \frac{\gamma^2}{\sigma^2}\right) \approx \frac{\gamma^2}{\sigma^2} \tag{9}$$

Substituting the above into eqn. 8 we obtain the following condition for reliable communication:

$$\frac{\gamma^2}{\sigma^2} > (2 \ln 2)\frac{J}{N} \tag{10}$$

where the N is the number of sites used to hide watermark information bits and J is the information rate. Eqns. 8 and 10 reduce to the more familiar form [1] if the 'bandwidth' B of the channel is set to half the number of sites, $N/2$. Note that the noise power can be considerably greater than the signal power and, in theory at least, the message may still be transmitted reliably!

The strategy for communicating the watermark is now clear. Because a watermark should be imperceptible the signal to noise ratio (SNR) is severely limited. Reliable communication can only be assured by increasing bandwidth B to compensate for poor SNR. Hence, in the case of watermarking the maximum number N of suitable transform domain coefficients should be exploited for hiding information in the image. An analogous situation occurs in satellite and mobile communications where SNR is limited by power restrictions at the transmitter. There are also many similarities to secret military communications where an opponent may also attempt to detect, intercept or block a transmission. Watermarking may be considered as being an application of *spread spectrum communications* [5].

The Shannon limit may be approached by applying error control codes. Robust error correction techniques can be employed if necessary. Methods for error control coding are described by Sweeney [28], Chambers [1] and Blahut [29].

Information theory also gives some insights into where the watermark should be placed. Let us assume that the image may be considered as a collection of parallel uncorrelated Gaussian channels which satisfy eqn. 3 above with the constraint that the total watermark energy is limited:

$$\sum_{i=1}^{N} \gamma_i^2 \leq E \tag{11}$$

Using eqn. 4 and assuming that the noise variances are not necessarily the same in each channel, Gallager [26] shows that the capacity is

$$C = \frac{1}{2} \sum_{i=1}^{N} \log_2\left(1 + \frac{\gamma_i^2}{\sigma_i^2}\right) \tag{12}$$

where σ_i^2 is the variance of the noise corrupting the watermark and γ_i^2 is the average power of the watermark signal in the ith channel. This is a more general form of eqn. 8. Capacity is achieved when

$$\gamma_i^2 + \sigma_i^2 = T_h \quad \text{if } \sigma_i^2 < T_h \tag{13}$$

$$\gamma_i^2 = 0 \quad \text{if } \sigma_i^2 \geq T_h \tag{14}$$

where the threshold T_h is chosen to maximise the sum on the left-hand side of eqn. 11 and thus maximise the

energy of the watermark. This result shows clearly that the watermark should be placed in those areas where the local noise variance σ_i^2 is smaller than threshold T_h and not at all in those areas where the local noise variance exceeds the threshold. Note that the simple analysis presented here assumes that the noise corruption suffered by the watermark, as a result of common forms of image processing, is Gaussian. This is not an accurate assumption to make in many cases. However, the Gaussian assumption is not a bad choice given that the aim is to derive rules and heuristics that apply *in general* to a number of fundamentally different different image processing scenarios. The Gaussian noise model leads to a tractable analysis in many cases. Theoretically, it can also be considered to be a general noise model because of its conservative nature. Three justifications for its adoption in the absence of any information regarding the noise statistics include the central limit theorem [30], Herschel's theorem [31] as well as the principle of maximum entropy [32, 33]. In addition, additive white Gaussian noise theoretically gives the most difficult conditions in which to attempt communication [26]. Hence, the Gaussian noise assumption is actually quite conservative. A full analysis of the channel based on accurate knowledge of the noise statistics would lead to more accurate values for the channel capacity but would also be complicated by the need to evaluate difficult multidimensional integrals.

6 Examples

Fig. 1 shows 'Lena' watermarked using bidirectional coding and blocks with borders [9]. The image is of size 512 × 512 pixels, the inner block size is 12 × 12 pixels and the pixels are incremented by 3 to transmit a binary '1' and decremented by 3 to convey a binary '0'. The mark is for all intents and purposes invisible in Fig. 1 but may be detected quite readily [24, 9] even after lossy compression and scanning have been carried

Fig. 1 *Lena weakly watermarked using bidirectional coding*

Fig. 3 *Lena watermarked using fourth-order Daubechies wavelets*

Fig. 2 *Lena strongly watermarked using bidirectional coding*

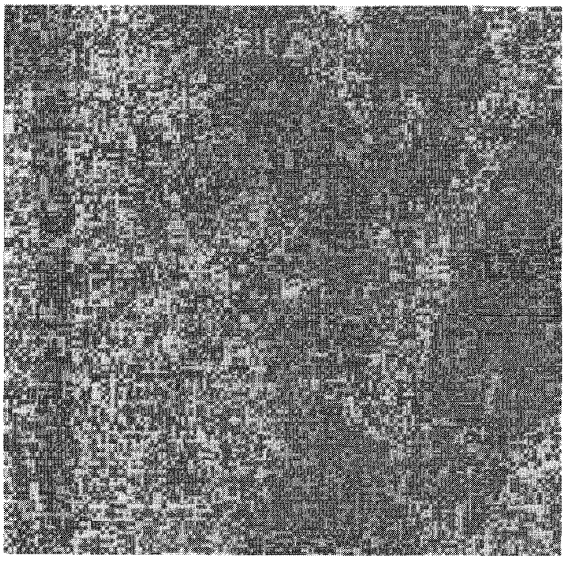

Fig. 4 *Watermark produced using Daubechies wavelets*

out. The watermark conveys 441 bits of information in ASCII form and the standard message reads: '012345 This is a watermark...'. Fig. 2 shows the same image strongly marked with a perturbation of ± 12 to make the mark readily visible.

Fig. 3 shows 'Lena' watermarked using the Daubechies wavelet transform. The block size is 8 × 8 and the maximum transform coefficient perturbation is ± 5. The watermark is conveyed by modulating 15 551 transform coefficients. The standard message is repeated a large number of times to occupy all of the available capacity. Note that the presence of the mark introduces no visible degradation.

Fig. 4 shows the difference between the wavelet marked version of the standard image and the original, scaled by a factor of 30 and offset by 127 grey scale levels. Fig. 5 shows a similar difference image for a watermark produced using the DCT. As in the case of the wavelet watermark, the DCT block size is 8 × 8 and the maximum transform coefficient perturbation is ± 5. The DCT watermark is conveyed by modulating 11 933 transform coefficients and the standard test message is repeatedly encoded as before.

Fig. 6 shows a watermarked image of a wolf on a snowy background. The image is of size 768 × 512. This image is very interesting from our point of view because it combines smooth background regions (the snow) with highly textured regions (the wolf). The watermark was produced using the Hadamard transform with an energy threshold $\varepsilon = 0.2$. The block size is 8 × 8 and the transform coefficient perturbation is ± 10. The watermark is conveyed by modulating 3840 transform coefficients. The absolute difference between the original image and the marked image, contrast enhanced using histogram equalisation, is shown in Fig. 7. In this case, areas with high information density (expressed in terms of the number of embedded watermark bits per block) are white, while areas which attract fewer watermark bits are darker. The outline of the wolf's head is quite clear. Note that, as before, information density is higher in textured regions.

Fig. 8 shows a segment of a watermarked image of Lena after JPEG [10] image compression followed by cropping. The size of the segment is 512 × 200 pixels. The watermark embedded in the uncropped image is 4096 bits long and the blocksize is 8 × 8 (i.e. just one bit per block). The encoded message consists of 32 bits ('0123' in ASCII). The watermark was placed using a DCT and the perturbation in the coefficient values was ± 10. JPEG was applied with a standard setting of 50 and no smoothing was used. By judicious use of concatenated error control codes [29, 1, 28] the watermark was recovered with ease from this cropped section.

It is apparent upon examining the watermarks in Figs. 4, 5 and 7 that the transform-based marking schemes possess a number of desirable features. One can mark according to the distribution of energy within the coefficients. In this way, one can place watermarks where they are least noticeable, such as within image texture and around edges. As a result, the watermark exhibits a ghost-like resemblance to the original image.

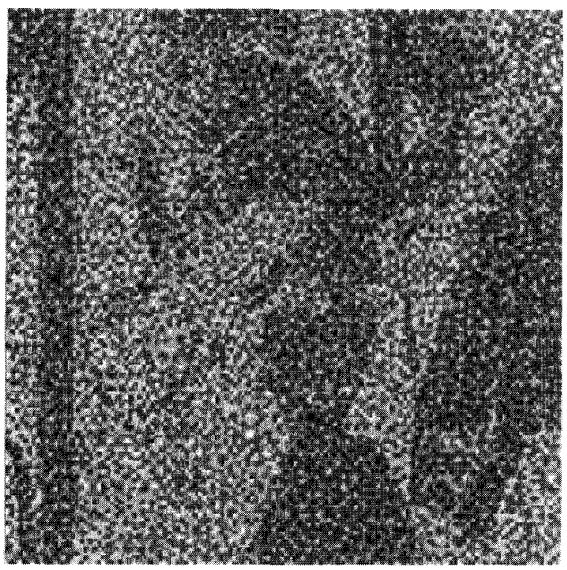

Fig.5 *Watermark produced using the discrete cosine transform*

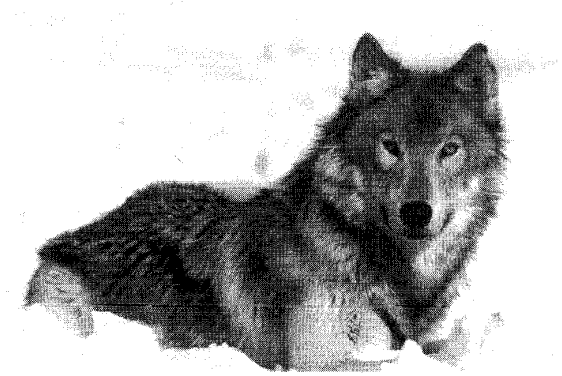

Fig.6 *Marked image of a wolf on a snowy background*

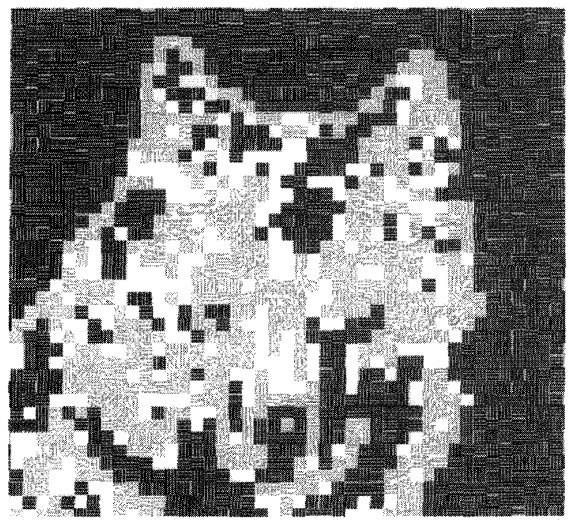

Fig.7 *Watermark around the region of the wolf's head*
The watermark was generated using the Hadamard transform. Areas with a high density of information are indicated by the brighter blocks

Fig.8 *Cropped grey scale image of Lena*
The size of image is 512 × 200 pixels

7 Conclusion

This paper has outlined a scheme for embedding robust watermarks in digital images. The watermarks are designed to be invisible, even to a careful observer, but contain sufficient information to identify both the origin and intended recipient of an image with a very low probability of error.

One key feature of the transform-based methods is that information bits can be placed adaptively, thereby making the watermark more robust to attack. A watermark is made imperceptible because it is designed to match the characteristics of the image to be protected. Transform-based methods have proven to be reasonably robust to image compression and standard image processing operations. In addition, transform-based methods yield a relatively large number of transform coefficients in which to embed the watermark. Future work will include the use of human visual models in designing watermarking schemes. The application of suitable error correction codes and digital signature techniques will also be investigated. In particular, the statistical characteristics of the *watermarking channel* need careful study. It is known that the distribution of the DCT coefficients of a typical image is well approximated by a Laplacian distribution [18]. It has been observed that the noise distortion imposed on the watermark by common image processing operations is non-Gaussian and impulsive in nature. Soft error control codes designed for additive wideband Gaussian noise (AWGN) channels (e.g. Reed–Muller codes) are not particularly effective in this application. The design of an optimal detector for the watermark depends on a clear knowledge of the noise statistics because such a detector can only be as good as the model assumptions upon which it is based. Finally, work will continue on devising watermarking schemes that do not require the original image to decode the watermark [21].

8 Acknowledgment

This work was supported by a Forbairt strategic research grant. The authors would like to thank Dr Peter J. Cullen for helpful advice and stimulating discussions.

9 References

1 CHAMBERS, W.G.: 'Basics of communications and coding' (Oxford Science Publications, Clarendon Press, Oxford, 1985)
2 HAYKIN, S.: 'Communications systems' (Wiley, 1994, 3rd edn.)
3 SCHNEIER, B.: 'Applied cryptography' (Wiley, 1995, 2nd edn.)
4 COX, I., KILLIAN, J., LEIGHTON, T., and SHAMOON, T.: 'Secure spread spectrum communication for multimedia'. Technical report, NEC Research Institute, 1995. ftp://ftp.nj.nec.com/pub/ingemar/papers/watermark.ps.Z
5 PICKHOLTZ, R.L., SCHILLING, D.L., and MILSTEIN, L.B.: 'Theory of spread spectrum communications-a tutorial', *IEEE Trans.*, 1982, **COM–30**, (5), pp. 855–884
6 BRASSIL, J., LOW, S., MAXEMCHUK, N., and O'GORMAN, L.: 'Electronic marking and identification techniques to discourage document copying'. Proceedings of INFOCOM 94, 1994
7 KURAK, C., and McHUGH, J.: 'A cautionary note on image downgrading'. Proceedings 8th Annual Computer Security Applications Conference, San Antonio, 1992
8 WALTON, S.: 'Image authentication for a slippery new age', *Dr. Dobb's J.*, 1995, **20**, (4), pp. 18–26, 82–87
9 DAUTZENBERG, C., and BOLAND, F.M.: 'Watermarking images'. Technical report, Department of Electronic and Electrical Engineering, Trinity College Dublin, 1994
10 PENNEBAKER, W.B., and MITCHELL, J.L.: 'JPEG still image compression standard' (Van Nostrand Reinhold, New York, 1993)
11 ZHAO, J., and KOCH, E.: 'Embedding robust labels into images for copyright protection'. Technical report, Fraunhofer Institute for Computer Graphics, Darmstadt, Germany, 1994
12 TIRKEL, A.Z., RANKIN, G.A., VAN SCHYNDEL, R.G., HO, W.J., MEE, N.R.A., and OSBORNE, C.F.: 'Electronic water mark'. Proceedings of Dicta–93, 1993, pp. 666–672
13 TIRKEL, A.Z., VAN SCHYNDEL, R.G., and OSBORNE, C.F.: 'A two-dimensional digital watermark'. Proceedings of ACCV, Singapore, 1995
14 VAN SCHYNDEL, R.G., TIRKEL, A.Z., and OSBORNE, C.F.: 'A digital watermark'. Proceedings of IEEE International Conference on *Image Processing*, Austin, Texas, 1994, pp. 86–90
15 VAN SCHYNDEL, R.G., TIRKEL, A.Z., and OSBORNE, C.F.: 'Towards a robust digital watermark'. Proceedings of Dicta-95, 1995
16 MATSUI, K., and TANAKA, K.: 'Video-steganography: how to secretly embed a signature in a picture'. IMA Intellectual Property Project Proceedings, January 1994, pp. 187–206
17 Ó RUANAIDH, J.J.K., DOWLING, W.J., and BOLAND, F.M.: 'Phase watermarking of images'. IEEE International Conference on *Image processing*, Lausanne, Switzerland, September 1996
18 CLARKE, R.J.: 'Transform coding of images' (Academic Press, London, 1985)
19 PRESS, W.H., TEUKOLSKY, S.A., VETTERLING, W.T., and FLANNERY, B.P.: 'Numerical recipes in C' (Cambridge University Press, 1992, 2nd edn.)
20 RAO, K.R., and YIP, P.: 'The discrete cosine transform: algorithms, advantages, applications' (Academic Press, 1990)
21 Ó RUANAIDH, J.J.K., BOLAND, F.M., and SINNEN, O.: 'Watermarking digital images for copyright protection'. Proceedings of the Electronic Imaging and Visual Arts Conference, Florence, February 1996. http://kalman.mee.tcd.ie/people/jjr/evapap.html.
22 LIM, J.S.: 'Two-dimensional signal and image processing' (Prentice–Hall International, 1990)
23 OPPENHEIM, A.V. and LIM J.S.: 'The importance of phase in signals', *Proc. IEEE*, 1981, **69**, (5), pp. 529–541
24 CARONNI, G.: 'Assuring ownership rights for digital images', *in* BRUEGGEMANN, H.H., and GERHARDT-HAECKL, W. (Eds): 'Reliable IT systems VIS '95' (Vieweg Publishing Company, Germany, 1995)
25 SHANNON, C.E.: 'A mathematical theory of communication', *Bell Syst. Tech. J.*, 1948, **27**, pp. 379–423, 623–656
26 GALLAGER, R.G.: 'Information theory and reliable communication' (Wiley, 1968)
27 PROAKIS, J.G.: 'Digital communication'. Series in electrical and computer engineering communications and signal processing (McGraw–Hill, 1995, 3rd edn.)
28 SWEENEY, P.: 'Error control coding: an introduction' (Prentice–Hall, 1991)
29 BLAHUT, R.E.: 'The theory and practice of error control codes' (Addison–Wesley, 1983)
30 PAPOULIS, A.: 'Probability, random variables and stochastic processes' (McGraw–Hill, 1984, 2nd edn.)
31 BRETTHORST, G.L.: 'Bayesian spectrum analysis and parameter estimation' (Springer–Verlag, 1989)
32 JAYNES, E.T.: '', *in* ROSENKRANTZ, R.D. (Ed.): 'Papers on probability, statistics and statistical physics, a reprint collection' (Kluwer, 1989)
33 BURTON, D., and FITZGERALD, W.J.: 'Bayesian parameter estimation: further results'. Technical report, Marconi Maritime Research Laboratory, Cambridge, England, 1989