

Data Hiding in Angle Quantisation

Ron G. van Schyndel ^{*1}

* Department of Physics, Monash University, Clayton 3168, Australia

ABSTRACT

In this paper, we present a non-linear method of embedding a signature in colour and monochrome images and demonstrate its recovery. The embedding process can be viewed as pseudo-random perturbations to angles between vector elements. The derived angles are dithered by the addition of a watermark, and encoded as a pseudo-noise sequence of dither angle offsets. This is followed by a re-quantisation for storage or transmission. The dither angles are recovered by scaling according to the pre-determined angle quantisation intervals. These intervals may be fixed according to some pattern, or they could be obtained adaptively from the local image. Performing a complex correlation with the known sequence enables recovery of sub-degree dither angles embedded in 8-bit data. This occurs without recourse to the original image.

This embedding process is additive in the angular domain and therefore multiplicative in the signal domain. Since the magnitude of the image vector is conserved, the image energy is largely unaltered by the embedding process.

Colour watermarks can be treated as sets of ordered triples (RGB), as pixel pairs in spatial or YIQ/YCbCr colour domain, or in a transform domain.

Keywords: Correlation, Angle Quantisation, Watermarking, Pseudo-noise Arrays, Complex Alphabets

1. INTRODUCTION

This paper describes results for an angle based encoding scheme that exploits discrete Cartesian to polar conversions to embed information onto arbitrary data types. This embedding process can be used in a data watermarking system, but can also be used as a diagnostic tool, or as a simple information carrier.

This research follows and extends an investigation of encoding watermarks in the hue angle of the HSV coded colour images¹⁶ to monochrome images, and by using other angle domains. Chae and Majunath¹ adopt a similar angle dithering strategy to that used here for the embedding of grey or colour images into the Y channel of YUV coded colour images. Watermarking in colour image data is also reported by Kutter³, and Tao and Orchard¹¹, where, for example, the blue channel is encrypted to minimise perceptual changes in the watermarked image.

The angles used here could be the hue angle directly or the angles could be obtained as the inverse tangent derived from pairs of pixel values found at a pre-determined pattern of image locations. These locations could themselves be pseudo-random, based on a key – adding another layer of protection.

The set of angles thus derived from the image is dithered by adding perturbing angles. The set of derived angles is pre-quantised, to permit the addition of dither to the derived angle vectors without overlap between adjacent angle bins.

The perturbing angles are typically downscaled values from some selected alphabet of angles, usually comprised of equal subdivisions over the range $0-2\pi$. The pattern of perturbing angles is chosen to be isomorphic to a spread-spectrum sequence or array of angles with near optimal auto-correlation properties^{9,13}. This helps maximise the watermark recovery process after estimation of the sequence of encoded dither angles.

If intended as a watermark⁶, further work is required to make this scheme more robust. The effects of JPEG image compression and small spatial distortions such as local warping and image cropping may be addressed as in Chae and

¹ Correspondence: email: Ron.vanSchyndel@sci.monash.edu.au, WWW: <http://www.physics.monash.edu.au/~ron>.
Ron van Schyndel is supported by an Australian Postgraduate Award.

Majunath¹, by encoding images after mapping to an appropriate transform domain using, for example, DCT or wavelet encoding.

Related angle encoding^{14,17} aimed to detect the digital delay times for the fundamental and overtones in discrete phase-angle encrypted carrier wave signals. The objective there was to characterise the properties of non-linear transmission channels rather than to watermark.

Section 2 outlines several schemes to derive angles from quantised pixel data and comments on the distribution of angles produced. Quantisation of these derived angles and the subsequent dither encoding mechanism is presented in Section 3. A typical “s-shaped” form for the decrease in correlative recovery results as the size of the embedded angles is decreased to make the watermark less perceptible. Section 4 analyses the contributing effects that cause the s-shaped loss of angle watermark recovery, and Section 5 discusses some strategies to improve the correlation further.

For simplicity, this paper describes the watermark embedding and recovery process using pseudo-noise (PN) sequences and correlation. It is understood that a more traditional angle-modulation scheme (such as MPSK) can also be used to embed information. This paper concentrates on the embedding process, not on the information coding.

2. OBTAINING ANGLES FROM PIXEL VALUES

We show here several methods of obtaining angles from an image, describing first the more general monochrome two-pixel method, and then its variations.

2.1 Two-pixel Vectors

We can define a derived complex image vector using $p_x, p_y \in \{0..1\}$ as

$$V = p_x + ip_y \tag{1}$$

by selecting, for example, pairs x and y of pixels, with values p_x and p_y . An angle can then be obtained from the ratio of their values as $\tan^{-1}(p_y / p_x)$. Adding a dither angle to the derived angle and then re-quantising the resulting total angle components, as p_x' and p_y' , perturbs each pair of original pixel values in a distributed, less perceptible, non-linear fashion.

The stored value of either or both selected pixels may change subtly with respect to the original values, for each encoded pixel entry. This reduces the dependence on and interaction with image content and greatly extends the available encoding space. It maintains, however, the same level of perceptible differences between original and watermarked images.

The relative position of the selected pixel pair is chosen arbitrarily. We chose to use a pattern of image coordinates generated by a pseudo-random sequence using a key to generate this sequence. This key can be used to encrypt the watermark and is required to decrypt it. Alternatively, the keys may be generated as in the DES algorithm, so that the watermark is asymmetric.

Each image pixel can only be dithered once when embedding a single watermark. Subsequent watermarks will interfere with the first.

This strategy has proven to be adequate for a range of natural images and for random test images with Gaussian or uniform distributions. Correlation effects between pixel pairs chosen randomly as above amounts to choosing two independent values from the same probability distribution (the image histogram).

Taking the *inverse tan* of integer ratios does place restrictions on the possible angles and all angles are not equally probable¹⁰. Some angles, such as 45°, are difficult to dither because of the large gaps to the next available angles with integer ratios.

The addition of a small dither to a vector needs to survive the re-quantisation of the pixel derived angle components from p_x to p_x' , so the length of the vectors associated with each pixel value pair also affects the dither encoding ability. The selected pixel pairs are filtered to ensure the vector length exceeds some minimum value. For images with a small number of quantisation levels, $0 < p_x, p_y < p_{max}$, the available angle space for encoding is severely restricted. Figure 1 shows the number of possible angles for square lattices with odd and even side-lengths p_{max} .

A ratio of positive integers produces angles over the range 0 to 90 degrees, with symmetry in the angle distribution around 45 degrees. Pixel values can be negatively offset, or interpreted as two's complement values, to produce a (related) 360° set, with 8-fold symmetry in the available image derived angles.

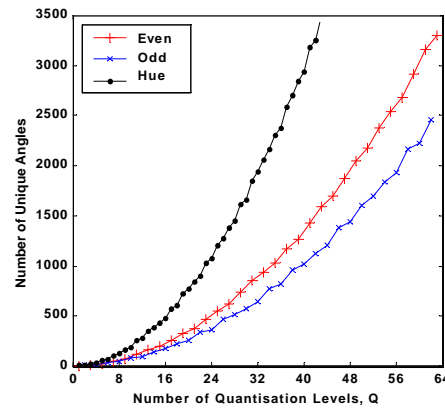


Figure 1. The number of possible angles constitutes a square lattice of even or odd side length Q , or a hex lattice of side length Q , representing the hue in HSV/HLS colour space.

2.2 One-Pixel Vectors

The angle generation strategy discussed above involving two pixel values can also be applied to single pixels as a variation of scalar quantisation.

In this case (1) becomes

$$V = p_x + i\sqrt{1 - p_x^2} \quad (2)$$

where $p_x \in \{0 \dots 1\}$. In other words, the image pixel is considered the real projection of a complex vector whose angle is dithered.

2.3 Colour Models

The angle q_o can be coded directly as the hue angle in the HSV or HLS colour space¹⁶. Since the hue angle is calculated using all three RGB components, the watermark is effectively spread over the three channels. Unfortunately, due to the non-linear mappings to obtain the image hue, separate processing of each RGB channel is likely to destroy the hue dither.

For two-channel colour descriptions such as YIQ or YCbCr, the angle q_o can be derived between the two colour channels in a manner analogous to the two-pixel vector model above. This mechanism suffers from the fact that there is usually a correlation between colour channels producing a non-uniform angle histogram – often concentrating around the troublesome 45°.

3. ANGLE QUANTISATION AND DITHER ENCODING

The set of derived angles is quantised to provide a fixed, minimum interval, of size $dq = 2p/s$, between all derived angle members (Figure 2). The gaps between the allowed quantised angle values are used to store the added dither angles, ensuring separability of the image values and the embedded watermark. The scale parameter s can range from 1 to 400 for 8-bit pixel data before dithered angles fail to be encoded and recovered sufficiently.




Figure 2. The four roots of unity for a sequence with an alphabet of 4 elements are compressed by the angle scale factor s and used to dither the original image angles.

The quantisation method provides superior isolation of watermark from image content. Other schemes use the mean or median of the local angle distribution to estimate the angle onto which the dither is added¹⁶. Here the local angle associated with any pixel is largely unrelated to the angle associated with the adjoining pixels, because of the use of a spatially disjoint pixel pair selection mechanism. This angle-encoded method is immune to local angle estimation techniques.

The dither angle to be embedded for each pixel pair corresponds to that of a unit vector with angles $2p/n$, where n is an integer. For binary sequences, the dither angle is effectively $\pm p/2$. When adding a binary dither angle to the

derived angle, the unit dither angle vector is added in a direction perpendicular to the derived angle, to maximise the separability of the watermark signal from the image derived angle.

The watermark dither angle can also be described in terms of a complex unit vector, $W_r = \exp(if_r)$, where f_r is the perturbation angle, which is applied to the original image vector V_o as a complex product

$$V_w = V_o W^{1/s} \quad (3)$$

The watermark is extracted by reconstructing and normalising the vector $V_w' = V_w / |V_w|$ from the quantised pixel values of the watermarked image, and subtracting the angle of the original normalised image vector, V_o' , if available.

$$f_e = s(q_w - q_o), \quad (4)$$

where $q_o = 0$, if the original image is unavailable.

The watermark estimate $W_e = \exp(if_e) = (V_w' / V_o')^s$, (where $V_o' = 1$ if the original image is not available), can then be cross-correlated with the reference watermark W_r .

Normalising the image vectors prior to cross-correlation constitutes a *phase-correlation*. Kuglin and Hines² first demonstrated the superiority of phase correlation over ordinary correlation for peak location and Thornton¹² extended this treatment to images.

Phase-correlating W_r with W_e produces

$$R_{er}(x, y) = \sum_{u=1}^N \sum_{v=1}^N W_e(u, v) W_r^*(u+x, v+y) \quad (5)$$




Figure 3. The sequence roots-of-unity are not aligned because the original q_o cannot be established from q_w . So b' and c' which should have aligned after multiplication by s , are now in opposition. If the original q_o is available, then the above vectors can be synchronised prior to multiplication by s , and the sequence roots-of-unity would then align.

Note that W_r and W_e are complex quantities and one may consider the a summation over an inner product of two vectors. W_r can thus be projected onto W_e and these projections summed.

$$|R_{er}(x, y)| = \sum_{u=1}^N \sum_{v=1}^N \cos(f_{er}) |W_e(u, v)| |W_r^*(u+x, v+y)| \quad (6)$$

where $f_{er} = |f_e(u, v) - f_r(u+x, v+y)|$. If a watermark exists, these projections will add constructively. W_r is by definition a unit vector, and by design W_e is also a unit vector, so the correlation between W_e and W_r becomes equivalent to

$$|R_{er}(x, y)| = \sum_{u=1}^N \sum_{v=1}^N \cos(f_e(u, v) - f_r(u+x, v+y)) \quad (7)$$

The effect of the correlation for different scales s can be seen by observing the argument of R_{er} as a histogram for each angle over the watermarked region of the image. So

$$|R_{er}(x, y)| = \sum_{\exists f_r} \sum_{f=-p}^{+p} P_{er}(f) \cos(f - f_r(x, y)) \quad \forall x, y \in W_r \quad (8)$$

represents the partial correlation for a particular angle in f_r and the total correlation is the summation over all different reference angles f_r . $P_{er}(f)$ represents a probability distribution of a particular angle of the estimated watermark.

3.1 Estimating the Original Image Angle, q_0

A problem with the watermark extraction model occurs when the original image angle q_0 is not available. Given the original image angle q_0 , a perfect recovery is possible up to the level of pixel quantisation errors using (4). If the original image is not available, a method of estimating q_0 must be found.

The up-scaling process of (4) by s also scales the q_0 component by s . This effectively randomises $(q_w - q_0)$ modulo 2π , and thus the recovered watermark is randomised. Figure 3 shows the problem arising from the rescaling of the perturbed data, where for example, components b' and c' now cancel.

One can achieve a more accurate estimate of q_0 by taking a spatial average of the local image angles (say, a 3×3 subregion). Langelaar et al⁴ have shown that the median filter is a good estimator to 'remove' an additive PN watermark in natural images. The angle watermark is multiplicative, but for small angles, it might be approximated by an additive watermark, so a spatial median filter of V_w might approximate V_o . Alternatively, since the watermark is additive in the angle domain, the spatial median filtering could be performed in this domain. Previously, we have shown¹⁶ some examples where spatial averaging was applied in the angle domain to estimate q_0 in a HSV encoded colour image.

This method works because of the local correlation that exists in a typical natural image if p_i and p_j are adjacent pixels. For an image of random-valued pixels or where p_i and p_j are far apart the local angle estimation method will not work at all since there is no spatial relationship between p_i and p_j .




Figure 4. A pixel-domain data-hiding scheme from [13], where the centre pixel value is estimated using its 8 neighbours. In the same way $N \cdot N$ (N odd) tiles of local neighbourhoods can be constructed. Since only the B pixels are watermarked, the watermark area will be $N^2/4$, and so a smaller SNR is expected. Similar processing could be carried out in the transform domain.

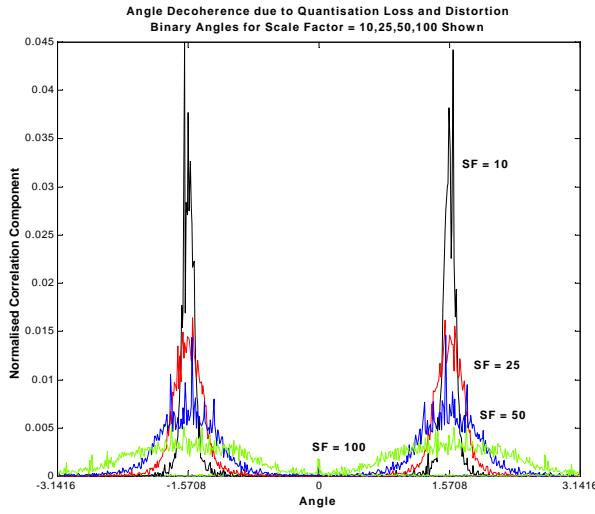


Figure 5. Histogram of recovered angles for a binary watermark. Curves for scale factors, SF of 10,25,50,100 are shown.

Mukherjee et al⁵ describes a method of estimating the original image values by imposing a dither modulation onto an image.

Mukherjee *et. al.* applied their method in an orthogonal transform domain, however they comment that the optimal predictors and estimators to be chosen for particular images may be hard to find.

Their methodology may be adapted to the spatial domain. In that case, the image pixels are subdivided into two classes: *A* and *B* (figure 4). Class *A* represents unperturbed image pixels and the values of class *A* are used to estimate the values of class *B*. These values of class *B* can be formed into angles and perturbed using our method. The perturbed or watermarked *B* values then *replace* the original *B* values. The unperturbed class *B* values can still be estimated from class *A*, which was not affected. They call the *B* class values, the *carrier* of the watermark.

Their method can be seen as a form of pixel location quantisation with the inter-quantile regions used as estimators for the original image values from which the watermark is determined.

Our method achieves the same result by pre-quantising the original image angle q_o ,

$$q_q = 2p/s \text{ int}(sq_o / 2p) \quad (9)$$

where *int* is integer rounding and q_o , q_q are expressed in radians. These quantised q_o are shown in Figure 2.

After scaling, sq_o will map to $2n\pi$, and the alignment of the watermark angles is assured. Effectively this synchronises the expanded perturbation angles on recovery (the zero angle location a' b' and c' in Figure 3 would always be aligned).

By effectively separating the angles occupied by the image from those occupied by the watermark, recovery has been substantially improved over the local-mean method. It is now also possible to recover the watermark where p_x and p_y are uncorrelated.

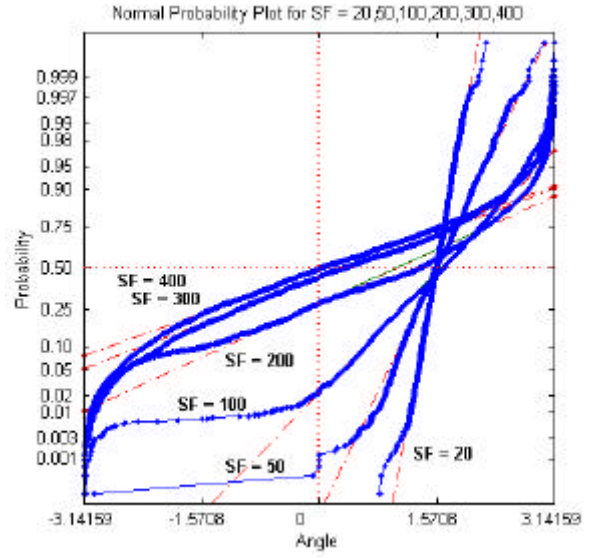


Figure 6. This graph shows a normal distribution fit to one side of the angle curves of Figure 5. The gradient is related to the standard deviation and the mean is the Y=0.5 intersection. It can be seen that 8-bit quantisation error will cause angles for scale factors above about 50 to wrap to the wrong quadrant. Increasing wrapping with higher factors eventually causes the mean to 'migrate' to 0, when total decorrelation occurs.

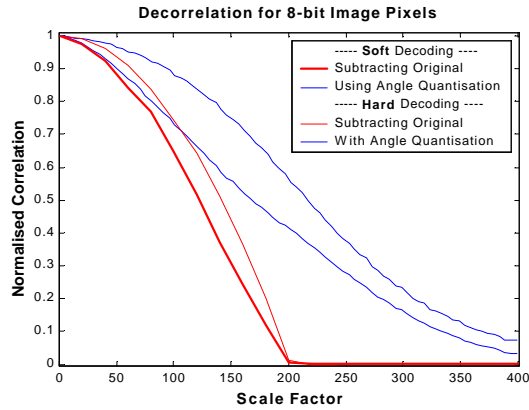


Figure 7. Relation between encoding scale-factor and degree of correlative recovery for a two-pixel vector and a uniformly distributed 8-bit noise image. This compares a non angle-quantised watermark with subtraction of original image, and an angle-quantised watermark without subtracting original. The scale factor is spatially constant. As shown, hard decoding of q_c to the nearest reference watermark angles improves the performance.

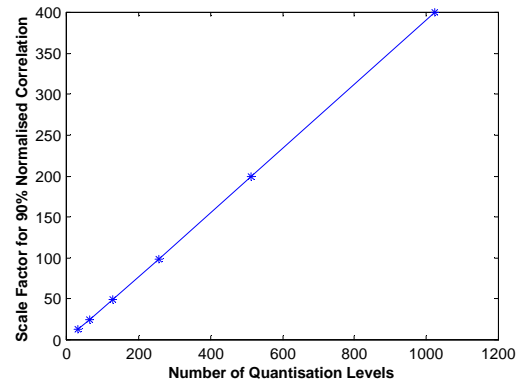


Figure 8. Dependence of the degree of correlative recovery with image bit depth (expressed as number of quantisation levels).

The near perfect separation of image and watermark also means that the watermark can be removed from the watermarked image given knowledge only of s . However, since s can be spatially varying (as $s(x,y)$) using a key-controlled mechanism, the watermark remains secure, if not robust.

4. MECHANISMS AFFECTING ANGLE RECOVERY

This section presents some of the results obtained for the dithered angle-watermarking scheme.

Figure 5 shows the histogram $P_{er}(q)$ of the recovered dither angles, for two reference angles corresponding to $\pm\pi/2$ and for four different scale factors, for the case of uniformly distributed 8-bit image data and a binary Legendre sequence. As the value of s increases, the two peaks in the recovered angles spread and eventually merge so that the dither vectors are misclassified. The widths of the peak in $P_e(f)$ can thus be related to the confidence of watermark recovery.

Figure 6 shows one side of the angle histogram of figure 5 in a different way. In this case, a Gaussian distribution would be plotted as a straight line with inverse gradient proportional to its variance, and the X intersection of the Y = 0.5 line as its mean. The graph shows that the mean for scale factors over 100 starts to drift towards 0. The departure of the curves from their fitted lines is due to angle wrap-around, which has already started at $s = 50$. Scale factors over 200 shows little change in their variance as the histogram now partially covers all recovered angles.

Figure 7 shows the loss of correlative recovery that occurs as the size of the scaling parameter s increases, making the added dither angles smaller. For 8-bit uniformly distributed data, the watermark can still be recovered for $s = 400$, where the added dither angle is about 0.2° , for the case of binary sequence encoding. At $s = 400$, the image signal to watermark ratio is about 50 dB in this case for a 251×251 watermark.

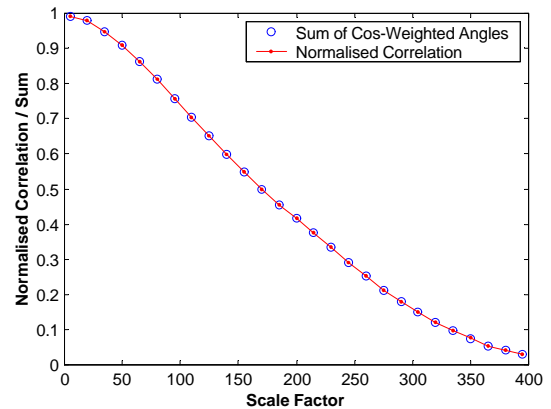


Figure 9. The sum of the recovered angles, weighted by the cosine of the distance to the template angle. This yields a distribution identical to the correlation curve of Figure 7.

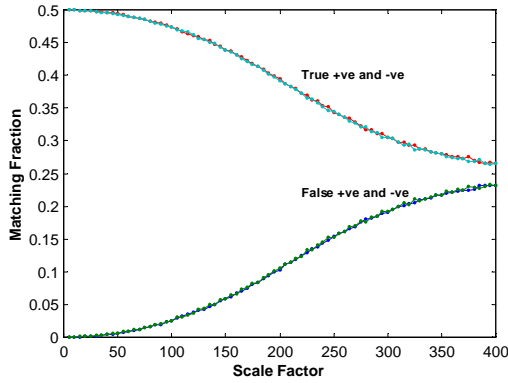


Figure 10. The fraction of recovered angles classified as true and false matches to the watermark template as a function of scale factor.

data and thus logarithmically with the number of quantisation levels. The Legendre sequence correlation value drops linearly with the number of sequence entries that are incorrectly encoded or recovered.

Figure 9 confirms that the degree of misalignment of the estimated dither vector angles accounts for the loss of correlation. The cosine of the estimated angle for each unit vector gives the useful component that adds constructively to the correlation. The sum of the cosine weighted unit vectors reproduces the s-shaped curve of Figure 7 with excellent precision. The sum of the sines of the recovered vector angles converges to zero.

Figure 10 shows the number of true and false classifications for the recovered angles of the unit vectors as a function of the scaling parameter s . The true and false curves converge when the dither angle estimation becomes random.

5. DISCUSSION

We can examine the distribution of V in (1) by creating a cooccurrence matrix with p_i and p_j as coordinates. Normally used in texture analysis to relate spatial patterns of gray values, when p_i and p_j are random and independent, the cooccurrence matrix becomes essentially a back-projection of the image histogram (pixel PDF) onto a 2D square array. The angle PDF is then the radial projection through this square array (see Figure 11).

If p_i and p_j are defined as Gaussian distributed with a non-zero mean, and $m = \sqrt{\sigma_i^2 + \sigma_j^2}$ is the mean radial distance of the Gaussian distributed p_i and p_j , from the origin, then the PDF f_q will be *Rician* distributed and

$$f_r(i, j) = \frac{r}{s^2} \exp\left(-\frac{r^2 + m^2}{2s^2}\right) I_0\left(\frac{rm}{s^2}\right) \quad (12)$$

$$f_q(i, j) = \frac{m \cos q}{s \sqrt{2p}} \exp\left(-\frac{m^2 \sin^2 q}{2s^2}\right) \left[1 - Q\left(\frac{m \cos q}{s}\right)\right] + \frac{\exp\left(-\frac{m^2}{2s^2}\right)}{2p}$$

where I_0 is the modified Bessel function of the first kind of order 0 and is defined as

$$I_0(x) = \frac{1}{2\pi} \int_0^{2\pi} \exp(x \cos q) dq \quad (13)$$

and Q is the Gaussian Tail Integral which is defined as

Figure 7 also compares the quantised angle watermark for random uniform data, with an un-quantised watermark where the original image is subtracted showing the improvement that angle pre-quantisation brings. The curves marked “Hard Decoding” are obtained by rounding the estimated dither angles to either $+\pi/2$ or $-\pi/2$ before correlation with the reference sequence. Using these “snapped” recovered angles increases the resulting complex correlation, as the estimated vectors all lie parallel or anti-parallel to the reference set. The improvement in correlation due to hard decoding decreases as the number of dither angles in the reference set is increased.

Figure 8 shows that the maximum value of s required to secure a fixed fraction (here 90%) of the possible peak correlation value scales linearly with the number of quantisation bits, V , used to represent the original image

$$Q(x) = \frac{1}{\sqrt{2p}} \int_x^\infty \exp\left(-\frac{y^2}{2}\right) dy \quad (14)$$

Figure 12 shows an image and its histogram (top left). By projecting a line radially from the origin of the cooccurrence matrix (bottom-left), the density is plotted as an angle histogram (bottom-right). It can be seen in this figure how two peaks can merge.

If we were to model the image histogram as a combination of N Gaussian peaks, we would require the sum of up to N^2 constructions of (12). Clearly, with such complex constructions there is a risk of overmodelling with even the simplest of image histogram.

One possible way to overcome the complexity of the image histogram is to transform the image to a domain whose angle histogram is more amenable to modelling.

Ramkumar⁸ describes such a sequence, that can be used to randomise an image and manipulate its PDF. He calls it the *Cyclic All-pass Sequence* and it is applied as a transform (the *CAP* transform) to our situation as follows.

Let $h \in R^N$ and $H = F(h)$, where $F()$ is the DFT, and let h be such that $|H(n)| = 1$ for $n = 0, 1, \dots, N-1$ so that $H.H^* = [1, 1, \dots, 1]$. Therefore

$$F^{-1}(H.H^*) = [1, 0, \dots, 0] \quad (15)$$

and since $F^{-1}(H.H^*)$ is the periodic auto-correlation of h , it follows that all h are mutually orthogonal.

Ramkumar continues that these unit vectors h , may have a randomly distributed phase, f_k , between $-\pi$ and $+\pi$ in even symmetry as follows:

$$f_k = \begin{cases} 0 \text{ or } p, & k = 0, \frac{N}{2} \\ q_k, & k = 0, 1, \dots, \frac{N}{2} - 1 \\ -q_{N-k}, & k = \frac{N}{2}, \dots, N-1 \end{cases} \quad (16)$$

and

$$H(k) = \cos(f_k) + i \sin(f_k), \quad k = 0, 1, \dots, N-1 \quad (17)$$

where $q_k, k = 1, \dots, N/2 - 1$ are randomly distributed between $-\pi$ and $+\pi$. In this case $h = F^{-1}(H)$ is a cyclic all-pass sequence.

A row of normalised image data, $p \in \{0, \dots, 1\}$ may then be transformed using this sequence as follows:

$$p' = F^{-1}(H.P), \quad (18)$$

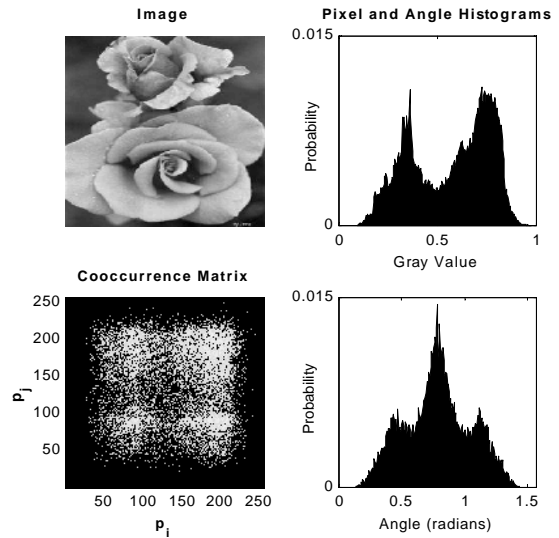


Figure 11. The image[†] (top-left) has a normalised image histogram with two peaks (top-right), which when projected onto a the rows and columns of the cooccurrence matrix (bottom-left) yields four peaks. The radial projection has merged two of the peaks to yield a normalised angle distribution with three peaks.

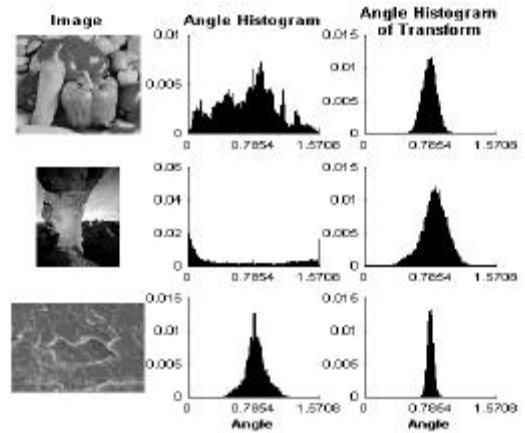


Figure 12. Three images[†], showing their widely different angle histograms but similar histograms after transformation using a 2D version of Ramkumar's cyclic all-pass filters.

where $P = F(p)$. p' is the transformed image, and $H.P$ is an element-wise product of H and P . Because of the even symmetry of H , p' will always be real, and because H is orthonormal, the process is reversed by:

$$p = F^{-1}(H^*.P'), \quad (19)$$

where $P' = F(p')$.

The histogram (or PDF) of p' is approximately Gaussian. If the random numbers that create H are derived from a 'seed' value, then this mechanism can be used as a simple watermark encryption, with the seed value as the symmetric encryption key. This means that the watermark is only extractable if the correct key is used.

The mechanism is readily extendable to two dimensions by casting H as a product array¹⁵, $H_{xy} = H_x \cdot H_y$, where H_x and H_y may be derived from different seed values.

Figure 12 shows three images, their angle histograms, $f(x,y)$, and their angle histograms after transformation, $y(x,y)$. The images were from F Petitcolas' image database⁷ and show those images with the most different $y(x,y)$.

These Gaussian-like histograms allows us to model most images as being Gaussian distributed. By transforming coordinates, $u = i-\eta$ and $v = j-\eta$, the Rician distribution of (12) becomes *Rayleigh* distributed.

The cooccurrence matrix of Figure 12 will then show a Gaussian-like distribution with only a changing width from image to image, and f_q is approximately uniformly distributed.

Figure 13 (left, top / bottom) shows the dramatic improvement to the peak height for a given scale-factor. This improvement comes at a price – the image PSNR (Figure 13 - right, top / bottom) is degraded as the CAP transform spreads artefacts from the angle perturbation over the whole image. Further work will entail whether a local watermark strength parameter can be incorporated into the CAP transform without compromising the desirable output histogram property.

In addition, Figure 13 (right, bottom, dashed line) shows the effect of angle quantisation for 1st quadrant angles near 0, $\pi/4$, and $\pi/2$. Effectively the loss in PSNR is due to quantisation of pixels near black or white, which predominate in this image (see its histogram) -- the image was chosen to highlight this effect.

CONCLUSION

We present a non-linear polar data embedding scheme which has potential applications in many areas. We have shown that a mechanism involving subtraction of the unembedded original signal also yields good results, but with angle pre-quantisation, it is possible to improve the recoverability of the embedded data.

That improvement may be increased if an image is transformed to an optimal domain (in this case one producing a Gaussian-like angle distribution).

ACKNOWLEDGEMENT

I would like to acknowledge the great help my supervisors, Dr Imants Svalbe and Dr. Andrew Tirkel have provided me over the years in my PhD study in this area.

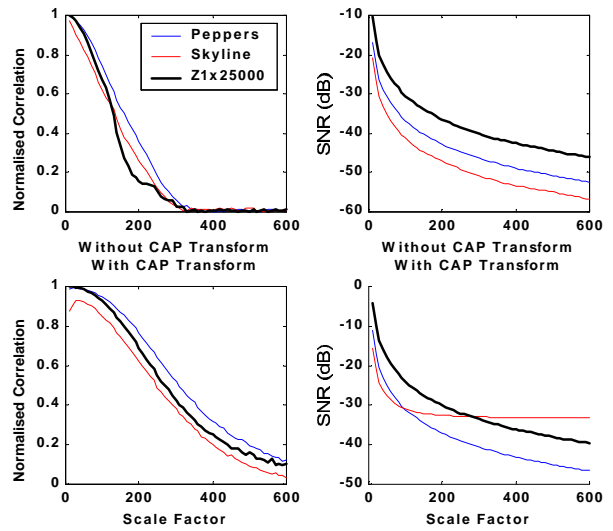


Figure 13. The more uniform decorrelation curves of the three images[†] of Figure 12 (left) after transformation using a 2D version of Ramkumar's cyclic all-pass filters. The corresponding signal-to-noise ratio (SNR) for the watermarks relative to the images are shown on the right.

REFERENCES

1. J. J. Chae and B. S. Majunath, "A Technique for Image Data Hiding and Reconstruction without Host Image", *SPIE Conference on Security and Watermarking of Multimedia Contents*, Vol 3657, San Jose, California, USA, January, 1999, pp. 386-396.
2. C. D. Kuglin, D. C. Hines Jr, "The Phase Correlation Image Alignment Method", *Proceedings of the International Conference on Cybernetics and Society*, pp. 163-165, 1975
3. M. Kutter, "Digital Signature of Color Images using Amplitude Modulation", *SPIE*, vol. 3022, San Jose, Feb., 1997, pp. 518-525.
4. G. Langelaar, R. Lagendijk, J. Biemond, "Removing Spatial Spread Spectrum Watermarks by Non-linear Filtering", *9th European Signal Processing Conference (EUSIPCO'98)*, September 1998, Island of Rhodes, Greece, pp. 2281-2284.
5. D. Mukherjee, J. J. Chae, S. K. Mitra, B. S. Manjunath, "A Source and Channel Coding Framework for Vector Based Data Hiding in Video", *IEEE Trans. Circuits and Systems for Video Technology*, 2000, Accepted.
6. J N. Nikolaidis, I. Pitas, "Digital Image Watermarking: an Overview", *IEEE International Conference on Multimedia Computing and Systems (ICMCS99)*, Volume 1, Florence, Italy, pp 1-6.
7. F. Petitcolas' image database² can be found at http://www.cl.cam.ac.uk/~fapp2/watermarking/benchmark/image_database.html as at Nov 2000
8. M. Ramkumar, *Data Hiding in Multimedia: Theory and Applications*, PhD Thesis, New Jersey Institute of Technology, 1999, Available at: http://njcmr.org/~mxr0096/thesis_abs.html as at Nov 2000
9. M. R. Schroeder, *Number Theory in Science and Communications*, 3rd Edn. 1997, Springer-Verlag.
10. I. D. Svalbe, "Natural Representations for the Hough Transform", *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 12, no. 2, 1991, pp. 336-342.
11. B. Tao and M. Orchard, "Coding and Modulation in Watermarking and Data Hiding", *SPIE Conference on Security and Watermarking of Multimedia Contents*, Vol 3657, San Jose, California, USA, January, 1999, pp. 503 – 510.
12. A. L. Thornton, *Colour object recognition using a complex colour representation and the frequency domain*, PhD Thesis, The University of Reading, Department of Engineering, May 1998
13. A. Z. Tirkel, R. van Schyndel, T. E. Hall, C. F. Osborne, "Secure Arrays for Digital Watermarking", *IEEE International Conference on Pattern Recognition (ICPR98)*, Brisbane, Australia, August 1998, pp. 1643-1645.
14. R. G. van Schyndel, A. Z. Tirkel, I. D. Svalbe, "Delay Recovery from a Non-Linear Polynomial Response System", *IEEE International Workshop on Intelligent Signal Processing and Communication Systems (ISPACS98)*, Melbourne, November 1998, Vol. 1, pp. 294-298.
15. R. G. van Schyndel, A. Z. Tirkel, I. D. Svalbe, T. E. Hall and C. F. Osborne, "Algebraic Construction of a new class of Quasi-orthogonal Arrays in Steganography", *SPIE Conference on Security and Watermarking of Multimedia Contents*, Vol 3657, San Jose, California, USA, January, 1999, pp. 354-364.
16. R. G. van Schyndel, A. Z. Tirkel, I. D. Svalbe, "A Multiplicative Color Watermark", *IEEE-EURASIP Workshop on Non-Linear Signal and Imaging Processing (NSIP99)*, Antalya, Turkey, 1999, pp. 336-340.
17. R. G. van Schyndel, "Using Phase-Modulated Probe Signals to Recover Delays from Higher-Order Non-linear Systems.", *IEEE Engineering in Medicine and Biology*, Melbourne, Australia, Accepted.

[†]The four images are as cited from Petitcolas' image database⁷: **Figure 12** - *Brandy rose*. Copyright photo courtesy of Toni Lanker, 17863 Hidden Acres Ln., Grand Haven, MI 49417 (jlanker@wmis.net). **Figure 13** - (a) *Peppers*. Courtesy of the Signal and Image Processing Institute at the University of Southern California, (b) *Skyline Arch*. Copyright photo courtesy of Robert E. Barber, Barber Nature Photography (BarberR@aol.com). (c) *Intergranular Stress Corrosion Cracking of an Al-Zn-Mg-Cu alloy*. Gérald Deshais, Department of Materials Science & Metallurgy, University of Cambridge.