Using Perceptual Models to Improve Fidelity and Provide Resistance to Valumetric Scaling for Quantization Index Modulation Watermarking

Qiao Li, *Student Member, IEEE*, and Ingemar J. Cox, *Fellow, IEEE*

***Abstract*—Traditional quantization index modulation (QIM) methods are based on a fixed quantization step size, which may lead to poor fidelity in some areas of the content. A more serious limitation of the original QIM algorithm is its sensitivity to valumetric changes (e.g., changes in amplitude). In this paper, we first propose using Watson's perceptual model to adaptively select the quantization step size based on the calculated perceptual "slack." Experimental results on 1000 images indicate improvements in fidelity as well as improved robustness in high-noise regimes. Watson's perceptual model is then modified such that the slacks scale linearly with valumetric scaling, thereby providing a QIM algorithm that is theoretically invariant to valumetric scaling. In practice, scaling can still result in errors due to cropping and roundoff that are an indirect effect of scaling. Two new algorithms are proposed—the first based on traditional QIM and the second based on rational dither modulation. A comparison with other methods demonstrates improved performance over other recently proposed valumetric-invariant QIM algorithms, with only small degradations in fidelity.**

***Index Terms*—Digital watermarking, perceptual model, quantization index modulation (QIM), rational dither modulation (RDM), valumetric scaling, Watson's model.**

I. Introduction

EARLY communications models of watermarking assumed a communication channel that contained two unknown noise sources. The first noise source modeled the interference due to the cover work and the second noise source modeled subsequent distortions between the time of embedding and detection. It is now recognized that digital watermarking can be best modeled as communication with side information [1]. The side information available is the original coverwork or host signal, which is entirely known to the watermark embedder. For such an arrangement, Costa showed [2] that the channel capacity of a communications channel with two noise sources—one of which is entirely known to the transmitter, but both unknown to the receiver, is equivalent to a channel in which the known noise source is absent. This result has important implications for digital watermarking—Costa's result implies that the coverwork (i.e., the known first noise source) need not interfere with the embedded watermark.

Quantization index modulation (QIM), first proposed by Chen and Wornell [3], provides a computationally efficient method for implementing codes based on Costa's work. QIM uses a structured lattice code to provide a computational efficient watermarking algorithm with high data capacity. However, the standard algorithm employs a fixed quantization step size which may lead to poor fidelity in some areas of the coverwork. Dither modulation (DM) has been introduced to partially addresses this issue. Section II describes QIM and DM in detail.

It is well known that improvements in fidelity and robustness can be achieved by adapting the watermark strength to the local perceptual characteristics of the coverwork. To address the fidelity issue, we apply Watson's perceptual model [4], described in Section III, to determine how much each discrete cosine transform (DCT) coefficient can be altered. This quantity, referred to as "slack," is then used to adaptively adjust the quantization step sizes used to quantize the DCT coefficients. The algorithm (DM–W) is described in Section IV. Experimental results on 1000 images indicate that the algorithm provides significant improvements in fidelity, as measured by Watson's distance, and performance degrades more gracefully with additive white Gaussian noise.

The most serious disadvantage of QIM has been its extreme sensitivity to valumetric scaling. Even small changes in the brightness of an image, or the volume of a song, can result in dramatic increases in the bit-error rate (BER). Several papers [5]–[10] have addressed this issue, and prior work is discussed in Section V.

The slacks computed by Watson's model do not scale linearly with valumetric scaling. As a result, our adaptive DM method [11] is not robust to valumetric scaling. In order to be robust to valumetric scaling, it is necessary that the quantization step sizes be scaled by the same scaling factor that the signal has undergone in order to correctly perform QIM decoding. A small modification to the Watson model leads to slacks that do scale linearly [11]. This algorithm, referred to as the modified Watson

Manuscript received January 3, 2006; revised February 1, 2007. This work was supported by the European Office of Aerospace Research and Development (EOARD) and in part by the Air Force Office of Scientific Research, Air Force Material Command, USAF, under Grant FA8655-03-1-3A46. The U.S. Government is authorized to reproduce and distribute reprints for Government purpose notwithstanding any copyright notation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the Air Force Office of Scientific Research or the U.S. Government. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Ton Kalker.

Q. Li is with the University College, London, Suffolk IP5 3RE, U.K.

I. J. Cox is with the Department of Computer Science and Department of Electronic and Electrical Engineering, University College London, Suffolk IP5 3RE (e-mail: ingemar@ieee.org).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TIFS.2007.897266

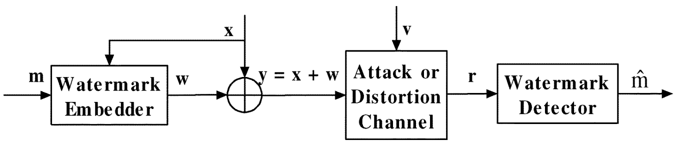

Fig. 1. Watermarking as a communication system.

model (DM–MW) is compared to the work of Oostveen *et al.* [7] and shown to have superior performance. The modification of the Watson model does lead to a degradation in fidelity. However, this degradation is slight, and the new algorithm still has significantly better fidelity when compared with either the original DM algorithm or that of Oostveen *et al.* [7]. This work is described in Section VI.

Dither modulation using a DM–MW implicitly assumes that the slacks calculated by the embedder and detector are the same, even though the function of embedding alters the DCT coefficients. Since these alterations are small, this assumption is usually true. However, this is also the source of some error. To guarantee that the slack is unaffected by the embedding procedure, we investigated using rational dither modulation.

Rational dither modulation (RDM) [12] has been recently proposed as an alternative DM method in which the quantization step size at time k is a function of the watermarked samples at earlier times. This algorithm is described in Section VII. If this function is chosen such that it scales linearly with amplitude, then RDM is invariant to valumetric scaling. To incorporate a perceptual model within the RDM framework, we propose to calculate the quantization step sizes for the current block k based on slacks of the previous watermarked block. Thus, while the slacks of the current block k are affected by the embedding process itself, the embedding of block k is based on the slacks from the previously watermarked block $k-1$, whose slacks were altered in the previous iteration but are unaffected by the processing of block k. Thus, both the watermark embedder and the watermark detector are guaranteed to use the identical values of step size (ignoring noise and roundoff error). This is described in Section VIII. Of course, once again, there is a degradation in fidelity, but experimental results indicate that it is small. Experiments shows that the final algorithm, rational dither modulation using modified Watson distance (RDM–MW), has the lowest BER as a function of amplitude scaling, compared to DM, RDM, DM–MW, and the algorithm of [7].

II. Quantization Index Modulation

Watermarking with side information is modeled by the communication system shown in Fig. 1. The message m and the cover Work or host signal x (i.e., image or song) are input into the watermark embedder, which outputs a watermark w that is added to the cover Work to produce the watermarked work y. The watermarked work then undergoes a number of distortions that are modeled as an unknown noise source v. The watermark detector receives a distorted, watermarked work r (i.e., $r = x + w + v$ and decodes a message $\hat{m}$).

A quantizer maps a value to the nearest point belonging to a class of predefined discontinuous points. The standard quantization operation with step size Δ is defined as

$$Q(x, \Delta) = \text{round}\left(\frac{x}{\Delta}\right)\Delta \tag{1}$$

where the function round(.) denotes rounding a value to the nearest integer.

Basic QIM uses two quantizers Q_1 and Q_2. They can be used to quantize the host signal to two sets of disjoint points, one set represents bit "0" while the other represents bit "1." When a message is being embedded, Q_1 or Q_2 is chosen according to the message bit to quantize x to the nearest quantization point. For example, Q_1 and Q_2 may be chosen such that Q_1 quantizes to even integers and Q_2 quantizes to odd integers. If we wish to embed a "0" bit, then Q_1 is chosen, else Q_2. The watermarked signal y is given by

$$y = Q_m(x, \Delta). \tag{2}$$

A. Dither Quantizer for QIM

Dither modulation is an extension of the original QIM algorithm, proposed by Chen and Wornell [13]. The purpose of the dither modulation, which is commonly used to deal with quantization noise [14], [15], is threefold. First, it is well known that a pseudorandom dither signal can reduce quantization artifacts to produce a perceptually superior quantized signal. Second, dither ensures that the quantization noise is independent of the host signal x. Third, the pseudorandom dither signal can be considered to act as a key which is only known to the embedder and detector, thereby improving the security of the system.

The host signal sample x_n is quantized with the resulting dithered quantizer to form the watermarked signal sample y_n. The embedding function embeds message bit m_n by

$$y_n(x_n, m_n) = Q(x_n + d(n, m_n), \Delta) - d(n, m_n) \tag{3}$$

where

$$d[n,1] = \begin{cases} d[n,0] + \frac{\Delta}{2}, & d[n,0] < 0, \\ d[n,0] - \frac{\Delta}{2}, & d[n,0] > 0, \end{cases} \quad n = 1, 2, \ldots, L \tag{4}$$

and $d(n, 0)$ is a pseudorandom signal usually chosen with a uniform distribution over $[-\Delta/2, \Delta/2]$, and L is the number of samples.

B. Hard- and Soft-Decision Detection

During detection, the detector calculates two signals $S_r(n, 0)$ and $S_r(n, 1)$ by embedding "0" and "1" into the received signal r separately, in the same manner as (3)

$$\begin{aligned} S_r(n, 0) &= Q(r_n + d[n, 0], \Delta) - d[n, 0]; \\ S_r(n, 1) &= Q(r_n + d[n, 1], \Delta) - d[n, 1]. \end{aligned} \tag{5}$$

The detected message bit is then determined by judging which of these two signals has the minimum Euclidean distance to the received signal r

$$\hat{m}_n = \underbrace{\text{argmin}}_{l \in 0,1} (r_n - S_r(n, l))^2. \tag{6}$$

The above description embedded one bit in each sample. In practice, we usually spread one message bit into a sequence of

N samples. One way to achieve this is to use a rate $1/N$ repetition encoding (i.e., simply embedding the same message bit in N samples). Detection can still be performed on a one-bit-per sample basis followed by a majority vote taken over the N samples to decide which message bit was embedded. We refer to this as hard-decision detection, which is given by

$$\hat{m_n} = \left\lfloor \frac{2}{N+1} \sum_{h=(n-1)N+1}^{nN} \underbrace{\operatorname{argmin}}_{l\in 0,1}(r_h - S_r(h,l))^2 \right\rfloor, \quad n = 1, 2, \ldots, \frac{L}{N}, \text{ where } \lfloor\ \rfloor \text{ is the floor function}. \tag{7}$$

An alternative detection strategy is to accumulate the two Euclidean distances for N samples and then determine the detected message bit, that is

$$\hat{m_n} = \underbrace{\operatorname{argmin}}_{l\in 0,1} \sum_{h=(n-1)N+1}^{nN} (r_h - S_r(h,l))^2, \quad n = 1, 2, \ldots, \frac{L}{N}. \tag{8}$$

The code rate is also $1/N$ but this soft-decision decoding [13] usually outperforms hard-decision decoding.

The original QIM and DM algorithms, described before, use a fixed quantization step size that is independent of the content. However, it is well known that the ability to perceive a change depends on the content. For example, the human visual system is much less sensitive to changes in heavily textured regions and more sensitive to changes in uniform regions. To account for this, we propose using a perceptual model to automatically select the quantization step size at each sample. Section III provides a description of the perceptual model we use, and Section IV then describes how this model is incorporated into the QIM framework.

III. Watson's Perceptual Model

Any perceptual model of the human visual system (HVS) has to account for a variety of perceptual phenomena, including luminance masking, contrast masking and sensitivity, all of which are discussed shortly. In psychophysical studies, the level of distortion that can be perceived in just more than 50% of experimental trials is often referred to as a just noticeable difference (JND). This difference is considered the minimum that is generally perceptible, and JNDs are sometimes employed as a unit for measuring the distance between two images.

Watson's model estimates the perceptibility of changes in individual terms of an image's block DCT.[1] It uses the block DCT transform, which proceeds by first dividing the image into disjoint 8×8 blocks of pixels. Each block is then transformed into the DCT domain, resulting in the block DCT coefficients of the transformed image C. We denote one term of the kth block by $C[i,j,k]$, $0 \leq i,j \leq 7$. $C[0,0,k]$ is the dc term (i.e., the mean pixel intensity in the block).

[1]Note that we are not referring to quantized JPEG coefficients.

Watson's model consists of a sensitivity function, two masking components based on luminance and contrast masking, and a pooling component.

A. Sensitivity

The model defines a frequency sensitivity table t. For each DCT coefficient (i,j), each table entry $t[i,j]$ is approximately the smallest discernible change in the absence of any masking noise. The resulting frequency sensitivity table is shown in [16]. Note that it is a table of constant values.

B. Luminance Masking

Luminance adaptation refers to the fact that a DCT coefficient can be changed by a larger amount before becoming perceptible, if the average intensity of the 8×8 block is brighter. The luminance-masked threshold $t_L[i,j,k]$ is given by

$$t_L[i,j,k] = t[i,j]\left(\frac{C_o[0,0,k]}{C_{0,0}}\right)^{\alpha_T} \tag{9}$$

where α_T is a constant with a suggested value of 0.649, $C_o[0,0,k]$ is the dc coefficient of the kth block in the original image, and $C_{0,0}$ is the average intensity of the image. Alternatively, $C_{0,0}$ may be set to a constant value representing the expected intensity of images.

C. Contrast Masking

Contrast masking (i.e., the reduction in visibility of a change in one frequency due to the energy present in that frequency) results in a masking threshold$s[i,j,k]$ given by

$$s[i,j,k] = \max\left(t_L[i,j,k], |C_o[i,j,k]|^{0.7} t_L[i,j,k]^{0.3}\right). \tag{10}$$

The final threshold $s[i,j,k]$ estimates the amounts by which individual terms of the block DCT may be changed before resulting in one JND. We refer to these thresholds as slacks.

D. Watson Distance

Given the original content c_o, its transform domain coefficients are denoted by C_o, and C_w denotes the watermarked transform coefficients. The Watson distance between the original and watermarked image is then given by

$$D_{\text{wat}}(c_o, c_w) = \left\{\sum_{i,j,k}\left(\frac{C_w[i,j,k] - C_o[i,j,k]}{s[i,j,k]}\right)^4\right\}^{1/4}. \tag{11}$$

Aside, we note that the JPEG standard provides a sample set of image-independent quantization matrices which, while not part of the standard, have been very widely adopted. The JPEG standard requires that the quantization matrix be part of the compressed file. Watson's model can be used to determine an image-dependent quantization matrix that can significantly improve the performance of JPEG. The interested reader is directed to [17] for further information.

It is interesting to contrast Watson's perceptual distance with measures of fidelity based on the document-to-watermark ratio

Fig. 2. Original image.

Fig. 3. Watermarked image (using DM method) with a Watson distance of 513, a $\mathrm{DWR} = 15$ dB, and a PSNR of 28 dB.

(DWR) or peak signal-to-noise ratio (PSNR). The DWR is defined as

$$\mathrm{DWR} = 10\log_{10}\left(\frac{\sigma_x^2}{\sigma_w^2}\right) \tag{12}$$

where x represents the original host signal, y is the watermarked signal, and the watermark $w = y - x$. Similarly, the PSNR is defined as

$$\mathrm{PSNR} = 10\log_{10}\frac{255^2}{\frac{1}{N}\sum_{n=1}^{N}(w_n)^2} \tag{13}$$

where N is the total number of pixels and $w_n = y_n - x_n$

Figs. 3 and 4 show two images with a $\mathrm{DWR} = 15$ dB and a PSNR of 28 dB. The original image is shown in Fig. 2. Clearly, Fig. 4 has a much higher fidelity than Fig. 3, yet the DWR and PSNR measures are identical. In contrast, the Watson distances for Figs. 4 and 3 are 53 and 513, respectively. Figs. 5 and 6 show the same images at a DWR of 35 dB and a PSNR of 49 dB. The Watson distances of Figs. 5 and 6 are 39 and 8, respectively. It is very difficult to discern any difference even viewed on a high-resolution computer monitor. However, if the images are magnified, there are subtle artifacts within the sky region.

Fig. 4. Watermarked image (using the DM–W method) with a Watson distance of 53, a $\mathrm{DWR} = 15$ dB, and a PSNR of 28 dB.

Fig. 5. Watermarked image (using DM method) with a Watson distance of 39, $\mathrm{DWR} = 35$ dB, and a PSNR of 49 dB.

Fig. 6. Watermarked image (using the DM-W method) with a Watson distance of 8, $\mathrm{DWR} = 35$ dB and a PSNR of 49 dB.

Finally, we note that the Watson distance is not normalized by the size of the image, while the DWR and PSNR measures are. Consequently, a Watson distance of say, 100, is substantially worse for small images than for large images.

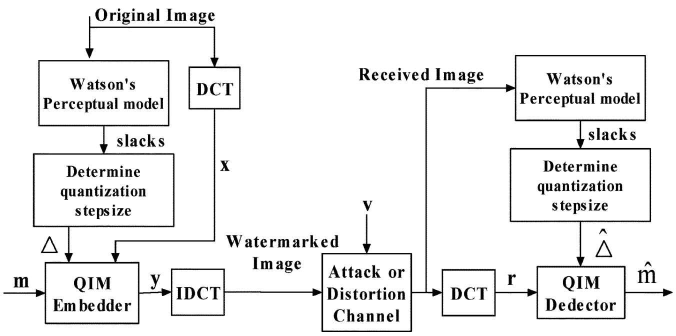

Fig. 7. Adaptive dither modulation based on Watson's model.

IV. Adaptive Dither Modulation Based on Watson's Perceptual Model

Each watermarking application may have its own specific requirements but the two most important and mutually conflicting requirements are usually fidelity and robustness. By locally adapting the quantization step size, we are able to provide significantly improved fidelity. For example, in regions of high texture, a larger step size can be used, while in regions of low texture, a small step size is chosen. As we shall demonstrate, this adaptivity can simultaneously improve robustness, at least in high-noise regimes.

DM is robust to additive white Gaussian noise, provided the standard deviation of the noise remains small compared with the quantization step size. For dither modulation, the distance between $d(n,0)$ and $d(n,1)$ is $\Delta/2$ (i.e., the distance between the embedding output for message bit "0" and "1" is $\Delta/2$). Thus, if the noise exceeds $\Delta/4$, the BER for DM rapidly degrades.

In contrast, an adaptive step size has, by definition, many different step sizes. Thus, we would expect the rate of change in BER to be slower. This is experimentally confirmed.

Traditional DM is nonadaptive, using a uniform scalar (step size Δ) for quantization as shown in (3). Note that the slacks, defined by (10), evaluate the amounts by which individual DCT coefficients may be changed according to Watson's perceptual model. This motivates us to design an adaptive DM method in which the DCT coefficients are quantized using step sizes that are based on Watson's perceptual model. We can use the slacks of (10) to adaptively select the quantization step size.

The adaptive DM system is schematically shown in Fig. 7.

The cover Work is converted to the DCT domain and the coefficients serve as the host signal x. The slacks from Watson's model are multiplied by a global constant G to determine the final quantization step size Δ for each DCT coefficient. The global constant G must be known to the detector and is the equivalent of the detector knowing the fixed quantization step size in traditional DM. The constant G is empirically adjusted to control the watermark strength and the document-to-watermark ratio. The message m is embedded by the DM embedder to obtain the watermarked signal y. After transmission, the received work r is used to estimate the corresponding quantization steps $\hat{\Delta}$. The estimation procedure is exactly the same as it is at the embedder. However, the estimate is now applied to the received, watermarked work rather than the original, unwatermarked work. Finally, using these step sizes, the message $\hat{m}$ is detected by the DM detector using (5) and (8).

Note that we use the original work to compute the quantization step size for each sample during embedding, and we use the distorted watermarked work to compute the quantization step size for each sample during detection. If these two step sizes are not the same, then a bit error may occur. In fact, even without distortions, there is the possibility that the slacks and, therefore, the quantization step sizes, computed at the detector, will be different due to the changes introduced by the embedded. However, in practice, very good correspondence is achieved, as is demonstrated below.

A. Experimental Results

We used a database of 1000 images from the Corel database, each of dimension 768×512. A binary message of length 12 288 bits is embedded into each image. We extracted 62 DCT coefficients from each 8×8 block, ignoring the dc and highest frequency coefficients. The entire sequence of 62×6144 coefficients was then pseudorandomized and each bit of the message was embedded in 31 random coefficients.[2] This is equivalent to embedding two bits in each block of the image.

To see the benefit that randomization of the host signal can bring, we examined the watermarking effectiveness of the embedder with and without randomization of the host coefficients. The effectiveness is simply the BER when detection is performed to the generated watermarked signal immediately after embedding (i.e., in the absence of any subsequent distortion). Given the experimental conditions described early in this section, it is observed that the BER or effectiveness with the nonrandomized method is 0.042. In contrast, the BER with randomization cross whole image is just 0.003. This demonstrates over an order of magnitude improvement with randomization.

Fig. 8 shows the BER as a function of additive white Gaussian noise. Results are provided for both our adaptive method and the original DM algorithm of [3]. In both cases, we adjusted the watermark strength such that the $\mathrm{DWR} = 35$ dB.

To fix the average DWR at 35 dB, we fixed the quantization step size at 3.0 for DM. For adaptive DM, the global constant G was adjusted to a value of 0.3 to meet the requisite DWR. Each point on a curve is the BER averaged over 1000 images from the Corel database.

The original DM algorithm has superior performance for low noise. We believe the poorer performance of the adaptive DM algorithm in the low noise region is due to a combination of 1) discrepancies between the corresponding estimated quantization step sizes at the embedder and the decoder and 2) the smaller step sizes used.

For the adaptive DM method based on Watson's model, the histogram of quantization step sizes used for each image is shown in Fig. 9. The overall average for 1000 images is 2.415, which is smaller than that used for standard DM. Note that the step sizes used for both methods quantize DCT coefficients, while the noise is added in the pixel domain of the image.

[2]Randomization of the coefficients ensures that the 31 coefficients associated with a single bit are 1) distributed spatially throughout the image and 2) distributed across a variety of low-, mid-, and high-frequency coefficients. This provides some robustness to 1) clipping and other spatially localized processing and 2) frequency filtering. The randomization also has the effect of whitening the host signal.

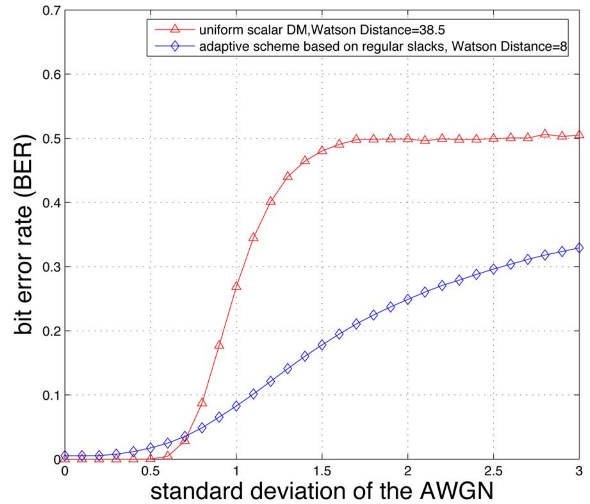

Fig. 8. BER as a function of additive white Gaussian noise ($\mathrm{DWR} = 35$ dB).

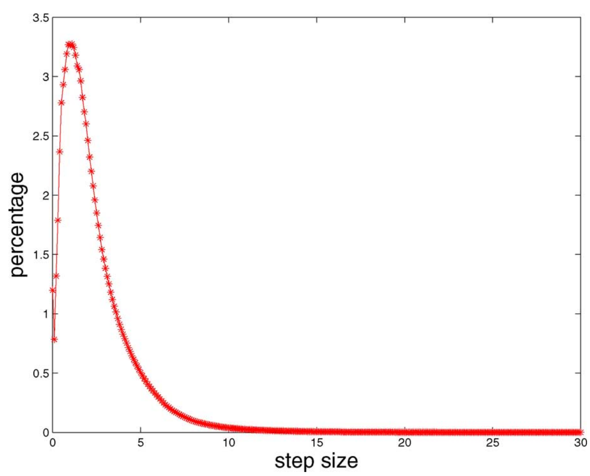

Fig. 9. Histogram of step size for adaptive DM using Watson's model. The average step size for 1000 image is 2.415.

When the standard deviation in the noise exceeds 0.7, which is roughly $\Delta/4$ ($\Delta = 3.0$), the BER for DM rapidly degrades, and the adaptive method is clearly superior. Note also that the superior performance of our algorithm is achieved with a very low Watson distance of 8 (i.e., very high fidelity) compared with the original method which has a Watson distance of 38.5. Thus, improved robustness and improved fidelity have been simultaneously achieved.

Unfortunately, despite the new method's superior performance under additive white Gaussian noise conditions, it remains vulnerable to amplitude scaling. When the amplitude of image is scaled by factor of β, the resulting luminance-masked threshold (denoted as $\widehat{t_L}[i,j,k]$) is calculated as

$$\widehat{t_L}[i,j,k] = t[i,j]\left(\frac{\beta C_o[i,j,k]}{\beta C_{0,0}}\right)^{\alpha_T} = t_L \quad (14)$$

that is, $\widehat{t_L}$ does not scale linearly with amplitude scaling, but is, in fact, invariant to amplitude scaling. Thus, referring to (10), the slack and corresponding quantization step size $\hat{\Delta}$ are not proportional to scaling factor β and the adaptive DM method is therefore sensitive to amplitude scaling.

V. Prior Work on Valumetric Robustness

There has been considerable work to develop QIM algorithms that are invariant to valumetric scaling. Eggers *et al.* [5] proposed to estimate the valumetric scaling by "securely embedd[ing] SCS pilot watermark." However, the fact that all watermarked content may contain the same pilot signal may lead to a security weakness—if the pilot signal can be estimated and removed, the watermark may not be detected. At the very least, the pilot signal is likely to reduce the watermark payload. Lee *et al.* [6] proposed estimating the global scaling factor using an EM algorithm, which does not need a pilot watermark. However, they note that the "complexity could be impractical." The closest work to ours is that of Oostveen *et al.* [7] which uses a simple perceptual model based on Weber's law. The quantization step size is a function of the average brightness of a neighborhood of pixels. This provides a simple perceptual model in which bright regions undergo larger changes than dark regions. It is obvious that if the image is scaled by a factor β then the quantization step size is scaled similarly.

Lagendijk *et al.* [9], [10] have presented several methods based on the characteristic function of the signal, and on a maximum likelihood procedure. Their algorithm requires models of both the host signal and noise. Experimental results are reported on synthetic data and real audio data. However, it is unclear how the algorithms will perform on real imagery.

Bas [8] recently proposed a method using so called "floating quantizers." The quantization and step sizes are based on the minimum and maximum of a triplet of pixels. A key advantage of this method is its robustness to nonlinear valumetric scaling such as gamma correction.

VI. Adaptive Dither Modulation Based on a Modified Watson Model

The previously described adaptive DM algorithm based on Watson's perceptual model remains sensitive to valumetric scaling, since the slacks do not scale linearly with amplitude scaling. To be robust to valumetric scaling, we need the slacks to scale linearly with valumetric scaling (i.e., we want the estimated $\hat{\Delta}$ to be multiplied by β when the amplitude of the signal is scaled by β). To this end, we modify the luminance masking in (9) to be t_L^M, given as

$$\begin{aligned} t_L^M[i,j,k] &= t_L[i,j,k]\left(\frac{C_{0,0}}{128}\right) \\ &= t[i,j]\left(\frac{C_o[0,0,k]}{C_{0,0}}\right)^{\alpha_T}\left(\frac{C_{0,0}}{128}\right). \end{aligned} \quad (15)$$

Compared with the previous (9), the last term is new. The term $C_{0,0}$ denotes the average of the dc components of the image, which we divide by 128 (the mean pixel brightness).[3] The modified slack is then given by

$$s^M[i,j,k] = \max\left(t_L^M[i,j,k], |C_o[i,j,k]|^{0.7} t_L^M[i,j,k]^{0.3}\right). \quad (16)$$

[3]Note that since the Watson distance and the proposed modified Watson distance are a function of $C_{0,0}$, the average brightness of the image, our methods will be susceptible to any cropping operation that changes the overall brightness value.

Thus, after the modification, when the image is amplitude scaled by factor of β, the resulting Luminance masking $\widehat{t_L^M}$ and slack $\widehat{S^M}$ are given by

$$\begin{aligned}
&\widehat{t_L^M}[i,j,k]\\
&=t_L[i,j,k]\left(\frac{\beta C_{0,0}}{128}\right)\\
&=t[i,j]\left(\frac{\beta C[0,0,k]}{\beta C_{0,0}}\right)^{\alpha_T}\left(\frac{\beta C_{0,0}}{128}\right)\\
&=\beta t_L^M[i,j,k]\\
&\widehat{s^M}[i,j,k]\\
&=\max\left(\widehat{t_L^M}[i,j,k],|(\beta C_o[i,j,k])|^{0.7}\widehat{t_L^M}[i,j,k]^{0.3}\right)\\
&=\max\left(\beta t_L^M[i,j,k],\beta^{0.7}|C_o[i,j,k]|^{0.7}\beta^{0.3}t_L^M[i,j,k]^{0.3}\right)\\
&=\beta\max\left(t_L^M[i,j,k],|C_o[i,j,k]|^{0.7}t_L^M[i,j,k]^{0.3}\right)\\
&=\beta s^M[i,j,k]. \qquad (17)
\end{aligned}$$

In the modified Watson model, the new luminance masking t_L^M and slack s^M scale linearly with β. The modified slack can then be used to determine the step size Δ_n^M

$$\Delta_n^M = G \times s_n^M, n = 1,2,\ldots,L. \qquad (18)$$

When the image is scaled by a factor of β, the estimated quantization step size $\widehat{\Delta_n^M}$ is also scaled by β. This provides an adaptive QIM algorithm that is theoretically invariant to valumetric scaling.

A. Experimental Results

For evaluation purposes, we again used a database of 1000 images, each of dimension 768×512. A binary message of length 12 288 bits was embedded into each image. In the following experiments, the global constant G was set to 0.075 for the modified Watson algorithm.

For comparison purposes, we evaluated the performance of the following algorithms:

A) the original nonadaptive QIM scheme of [13] using hard-decision detection;
B) the adaptive QIM scheme of [7] using the hard-decision detection;
C) adaptive QIM based on regular Watson model, soft-decision detection;
D) adaptive QIM based on the modified Watson model, hard-decision detection;
E) adaptive QIM based on the modified Watson model, soft-decision detection.

We compared the performance of these methods to amplitude scaling when the DWR is fixed at 35 and 25 dB. However, while the DWR is the same for images watermarked with the five algorithms, the average Watson distance between the watermarked and original image differs considerably, as shown in Table I.

Table I shows that the three adaptive schemes proposed here have very much lower perceptual distortion as measured by Watson's distance. Importantly, the modification to the Watson distance used in methods D and E to provide robustness against valumetric scaling produces only a small degradation in image quality and remain much better than methods A or B.

TABLE I
AVERAGE WATSON DISTANCE FOR DIFFERENT METHODS FOR A DWR OF 35 dB

Scheme	Watson Distance
(A) DM	38.5
(B) Oostveen	55.2
(C) DM-W	8
(D) DM-MW(HD)	9.4
(E) DM-MW	9.4

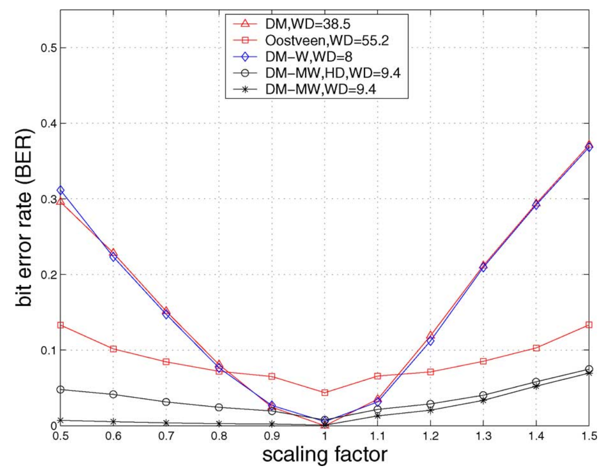

Fig. 10. BER versus amplitude scaling (DWR = 35 dB).

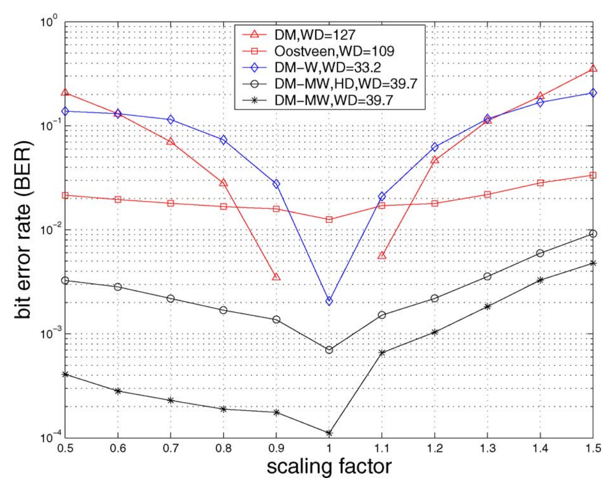

Fig. 11. BER versus amplitude scaling ($\mathrm{DWR} = 25$ dB). Note that for basic DM, the BER when there is no (unity) scaling is 0, and this point is therefore not plotted.

The robustness to amplitude scaling for all schemes is shown in Figs. 10 and 11 for DWRs of 35 and 25 dB, respectively. The performance is qualitatively similar for both DWRs though, of course, the BER is considerably smaller for $\mathrm{DWR} = 25$ dB, since the watermark is stronger in this case. The discussion below is therefore restricted to the case of $\mathrm{DWR} = 35$ dB.

We observe that for very small changes in scale $0.9 \le \beta \le 1.1$, the original algorithm A performs as well or better than the others. Our method C has poorer performance in this range, but for larger scale changes, it has similar or superior performance.

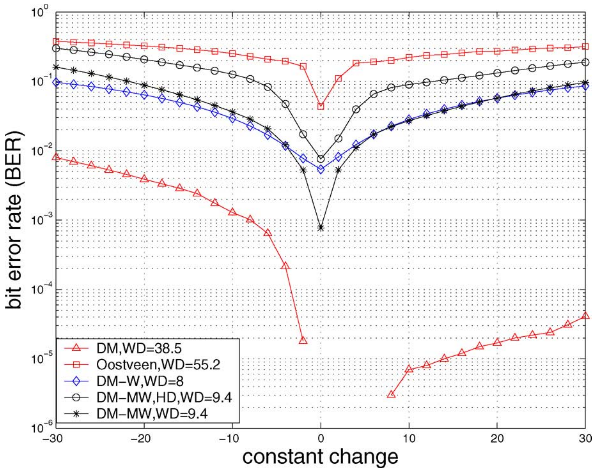

Fig. 12. BER as a function of constant luminance change ($\mathrm{DWR} = 35$ dB). Note that for basic DM, the BER when there is no change in luminance is 0, and this point is therefore not plotted.

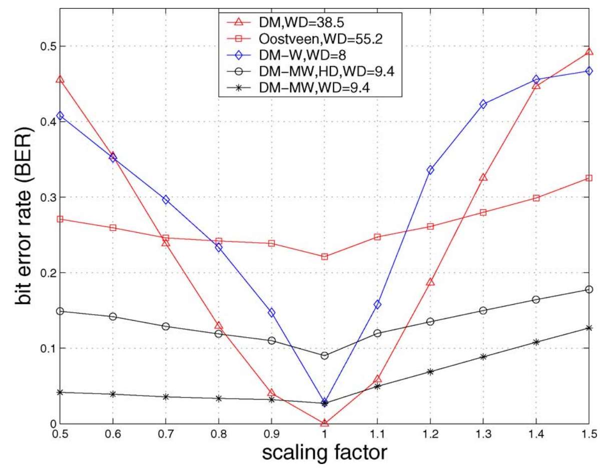

Fig. 13. BER as a function of amplitude scaling after a constant luminance change of 10 ($\mathrm{DWR} = 35$ dB).

It is also important to note that this is achieved with a perceptual distortion, as measured by Watson distance, of less than 20% compared with method A (see Table I).

Both algorithms A and C are not designed to be invariant to valumetric scaling. Bit-error rates of greater than 10% occur for $\beta < 0.8$ and $\beta > 1.1$. In contrast, Oostveen *et al.*'s method B and our methods D and E show much better robustness to scale changes. Clearly, method E outperforms all others with a BER that never exceeds 7% over the range of β tested. To ensure that this performance was not due to soft decoding alone, we implemented method D, DM-MW with a hard decision. While performing worse than method E, method D is still superior to Oostveen *et al.*'s method (which also uses hard decoding). One possible explanation for the poorer performance of Oostveen *et al.*'s method is that it is pixel-based rather than DCT based. Consequently, clipping of the pixels at 0 and 255 may have a more direct and deleterious effect.

Finally, we again note that while the perceptual distortion introduced by methods D and E is greater than for method C, the modification to the Watson model has only resulted in a small degradation in quality. Importantly, methods D and E have considerably higher quality than previous algorithms A and B.

Fig. 12 illustrates the sensitivity of all five algorithms to the addition/subtraction of a constant luminance value.

The method of Oostveen *et al.* is sensitive to this procedure because their adaptive step size is based on the average intensity of pixels in a neighborhood. This is clearly demonstrated in Fig. 12, where it has the highest BER. Conversely, the regular DM algorithm performs best, as the quantization occurs in the DCT domain and does not include the dc coefficient. Thus, we expect that DM would not be affected by this operation. Nevertheless, performance degradation is observed, particularly as we darken (subtract) the image. We suspect that this is due to clipping artifacts. Our proposed algorithms are substantially worse than DM, but about an order of magnitude less sensitive compared with Oostveen *et al.* The sensitivity of our methods is due to the fact that the slack calculations are based on the average intensity of the image, see (16), which is altered by the addition of a constant luminance.

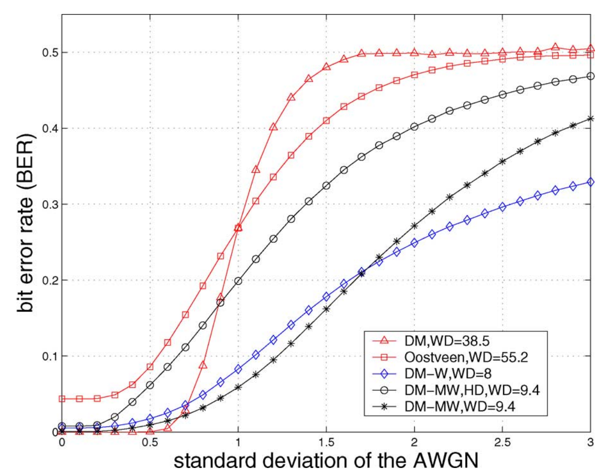

Fig. 14. BER as a function of additive white Gaussian noise (DWR = 35 dB).

Fig. 13 shows the BER as a function of amplitude scaling after the addition of a constant luminance change of 10. The performance of all algorithms is qualitatively similar. However, the BER is, on average, considerably worse than for amplitude scaling alone, as shown in Fig. 10.

Fig. 14 shows the sensitivity of all five algorithms to additive white Gaussian noise. The two curves from Fig. 8 are included for completeness. All of the adaptive methods perform worse than regular DM, for low levels of noise. However, the adaptive methods degrade more gradually as the standard deviation of the noise increases above about 0.6. The best performing algorithms in this high-noise regime are DM–MW using soft decoding and DM–W. Interestingly, the DM–W demonstrates the best performance for noise with standard deviations of greater than 1.5.

To understand why the performance of DM–W outperforms DM–MW for the noise of standard deviation greater than about 1.5, we examined the cumulative distribution in quantization step sizes for the two algorithms, as shown in Fig. 15. The two curves diverge at a step size of about 2. Thereafter, the DM–MW

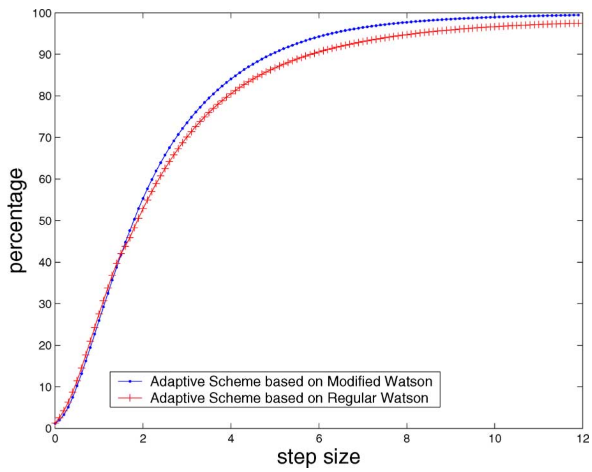

Fig. 15. Cumulative histogram of quantization step sizes for DM–WM and DM–M, (measured over all 1000 images).

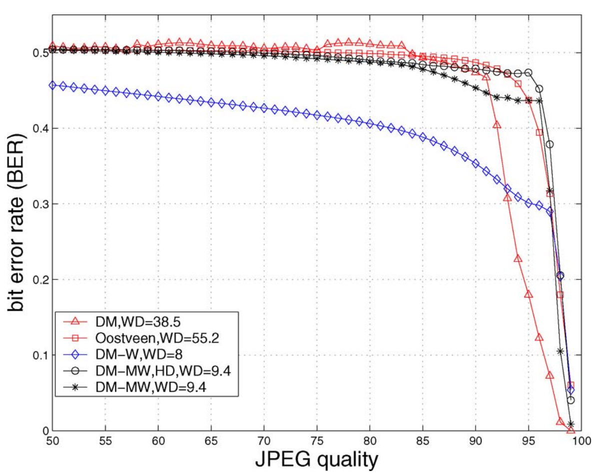

Fig. 16. BER as a function of JPEG quality for a fixed DWR of 35 dB.

increases more quickly, indicating that the DM–MW has smaller quantization step sizes than DM–W.

It should be noted that this performance difference is specific to a particular set of images. This is because the change in step size is due to the last term of (15). For bright images, this term will be greater than one and lead to larger step sizes than for the standard Watson model. Conversely, for dark images, this term will be less than one and lead to step sizes that are smaller than the standard Watson model. Thus, a different image test set might show that the DM–MW outperforms DM–W.

Fig. 16 shows the sensitivity of all five algorithms to JPEG compression for a fixed DWR of 35 dB. All of the algorithms are extremely sensitive to JPEG compression, although the original DM is somewhat more robust. However, for quality factors of less than 92%, dither modulation with Watson's model (DM–W) performs best. It is unclear why this should be.

Fig. 17 examines the sensitivity of all five algorithms when the Watson distance is fixed at 55. That is, we maintain a fixed fidelity rather than a fixed DWR. In this case, we see that DM and the adaptive method of Oostveen *et al.* perform significantly worse that our three adaptive methods. Once again, DM–W has the best performance for quality factors that are less than 75.

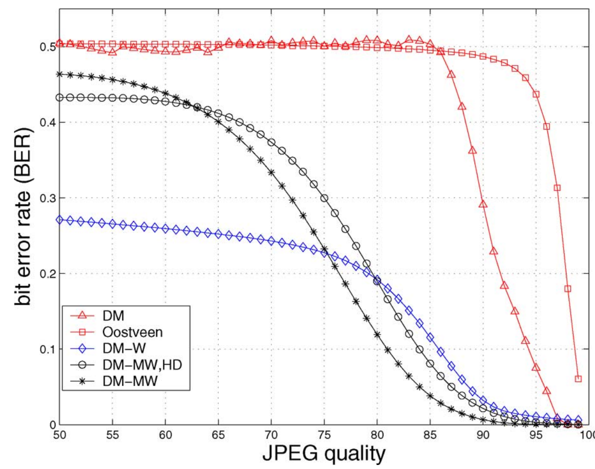

Fig. 17. BER as a function of JPEG quality for a fixed Watson distance of 55.

VII. Rational Dither Modulation

An investigation as to the source of errors in algorithm E, dither modulation using a modified Watson model and soft decoding, revealed that one important source of residual errors was due to the difference in slacks calculated at the embedder and detector. To eliminate this problem, we investigated combining our perceptual model with rational dither modulation.

RDM was first proposed by Perez-Gonzalez *et al.* [12] and is intended to provide valumetric invariance to QIM. Given a host signal $x = (x_1 \ldots x_N)$ and a watermarked signal $y = (y_1 \ldots y_N)$, then the kth bit of a watermark message $m = (m_1 \ldots m_M)$ is embedded as

$$y_k = g\left(y_{k-L}^{k-1}\right) Q_{m_k}\left(\frac{x_k}{g(y_{k-L}^{k-1})}\right) \tag{19}$$

where y_{k-L}^{k-1} denotes the set of past watermarked signals $(y_{k-L} \ldots y_{k-1})$ and the function $g()$ maps its L-dimensional input vector to a real value and has the property that for any valumetric scaling factor $\beta > 0$

$$g(\beta y) = \beta g(y). \tag{20}$$

This definition of RDM is intrinsically invariant to valumetric scaling. The function $g()$ can be chosen from a very large set of possible functions, including the L_p-norms, that is

$$g\left(y_{k-L}^{k-1}\right) = \left(\frac{1}{L}\sum_{i=1}^{L} |y_{k-i}|^p\right)^{1/p}. \tag{21}$$

However, it is well known that L_p-norms poorly model the human perceptual system. Thus, it is interesting to consider a function $g()$ that models properties of perception and satisfies the constraint of (20). An obvious candidate function is the slack function of (16) using the modified Watson model.

VIII. RDM With Modified Watson's Model

In Section VI, we used the modified slack to adaptively set the quantization step size and thereby provide robustness to valumetric scaling. Fig. 7 is a block diagram of the system. Notice that the quantization step size is determined by a local neighborhood around the host sample, x_k, and that this neighborhood is altered during the embedding of the watermark. Thus, during detection, we must rely on the fact that these alterations are small, and hope that the slacks based on the modified local neighborhood are the same as those determined during embedding. While this is often true, rational dither modulation suggests an alternative approach, in which the perceptual slack at time k is based on a nearby neighborhood of previously watermarked samples. Clearly, there may be some degradation in perceptual quality since a perceptual estimate made in a nearby neighborhood is not guaranteed to be perceptually relevant.

A. Implementation of RDM–MW

Before describing the implementation of rational dither modulation with a modified Watson perceptual model, we first describe an implementation of RDM in the DCT domain. This algorithm, denoted as RDM, is used for comparison.

Our implementation of RDM quantizes the 62 DCT coefficients of each 8×8 block (excluding the dc and highest frequency terms). For each DCT coefficient, we use its corresponding DCT coefficient from the previously watermarked 8×8 neighboring block to determine the quantization step size. Thus, the window size is 1. The function $g()$ is chosen to be the absolute value of the DCT coefficient (i.e., we use an L_1-norm in (21)) scaled by a global constant that is chosen so that the DWR averaged over all watermarked images is equal to a desired value.

We believe that a window size of 1 provides the fairest comparison with our adaptive RDM algorithm. However, we also implemented an RDM algorithm with a window size of 62, denoted RDM-62-L2-Norm, that uses an L2-norm of the 62 DCT coefficients in the previous block.

Our perceptually adaptive RDM method is denoted RDM–MW. To incorporate a perceptual model within the RDM framework, we propose to calculate the quantization step sizes for the current block k based on the slacks of the previous watermarked block. That is, each DCT coefficient in the current block is quantized based on the slack of its corresponding coefficient in the previously watermarked block. Thus, this method also has a window size of 1.

The slacks used to quantize block k are unaffected by the embedding process since they are determined from the previously watermarked block. Thus, both the watermark embedder and the watermark detector are guaranteed to use the identical values of quantization step size (ignoring noise and roundoff error). Of course, we expect a degradation in fidelity since we are basing our perceptual model of block k on calculations performed on block $k-1$. However, provided there is sufficient spatial continuity in the image, then this degradation is likely to be small. This is supported by the experimental results.

For all methods, a message of length 8192 is embedded using a 1/31 rate repetition code (i.e., one message bit is embedded in 31 DCT coefficients). As noted previously, randomization of the DCT coefficients significantly improves performance. However, randomization of the DCT coefficients is problematic for RDM. This is because the adaptive step size depends on the previous block of coefficients. Therefore, the previous neighboring block of coefficients must be watermarked prior to the current block (i.e., not in random order). In [18], we proposed a solution to this problem based on partitioning the image into disjoint regions, selecting a random block from each region, and only randomizing the coefficients in these blocks. While very satisfactory results were obtained, it was pointed out that a simpler solution is to randomize the message code, prior to embedding. Thus, we randomize the 8192×31 length repetition code and sequentially embed it in the DCT coefficients.

TABLE II
AVERAGE WATSON DISTANCE FOR VARIOUS METHODS FOR A DWR OF 35 dB

Scheme	Watson Distance
DM-MW	9.4
RDM-MW	10.2
RDM	54
RDM-62-L2-Norm	29

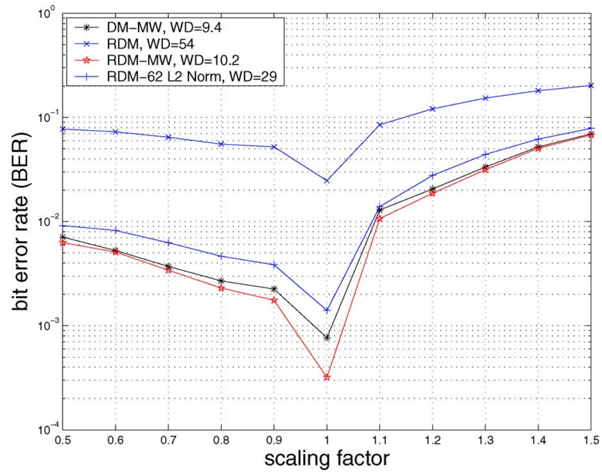

Fig. 18. BER as a function of valumetric scaling ($\mathrm{DWR} = 35$ dB).

IX. EXPERIMENTAL RESULTS

Once again, we watermarked 1000 images from the Corel database. We embedded the equivalent of 2 bits in each 8×8 block (i.e., 31 coefficients per bit). In all experiments, the average DWR was fixed at 35 dB.

The average Watson distance of the watermarked images for each of the three RDM algorithms and DM–MW is tabulated in Table II. As expected, the perceptual distortion for RDM–MW is worse than for DM–MW. However, the difference is small. RDM has a much larger Watson distance of 54 and RDM-62-L2-Norm has a perceptual distance of 29. Clearly, RDM–MW has provided a much reduced perceptual distortion.

The BERs as a function of valumetric scaling are shown in Fig. 18. The RDM–MW has the lowest BER. Surprisingly, our implementation of RDM with a window size of 1 has considerably worse performance.

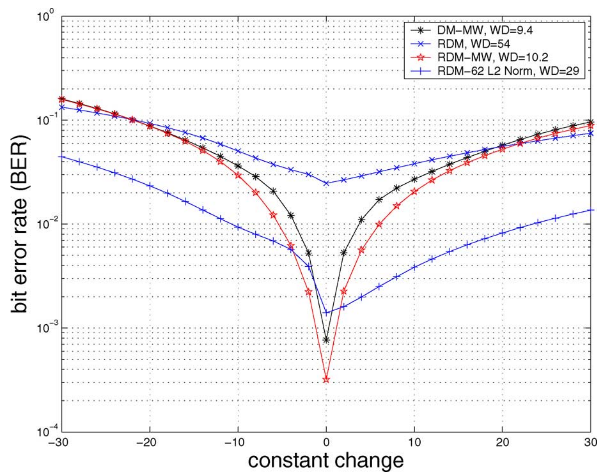

Fig. 19. BER as a function of constant luminance change ($\mathrm{DWR} = 35$ dB).

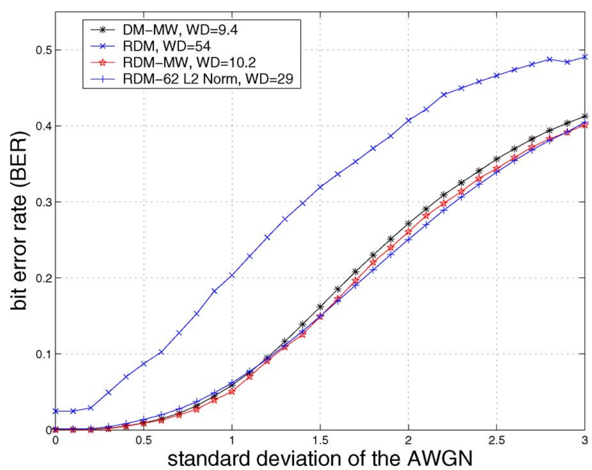

Fig. 20. BER as a function of additive white Gaussian noise (DWR = 35 dB).

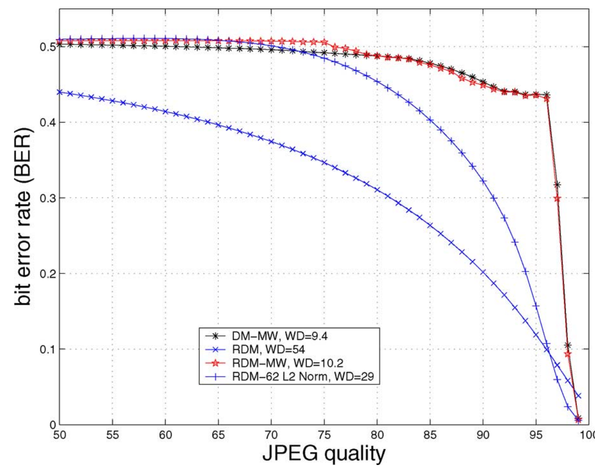

Fig. 21. BER as a function of JPEG quality for a fixed DWR of 35 dB.

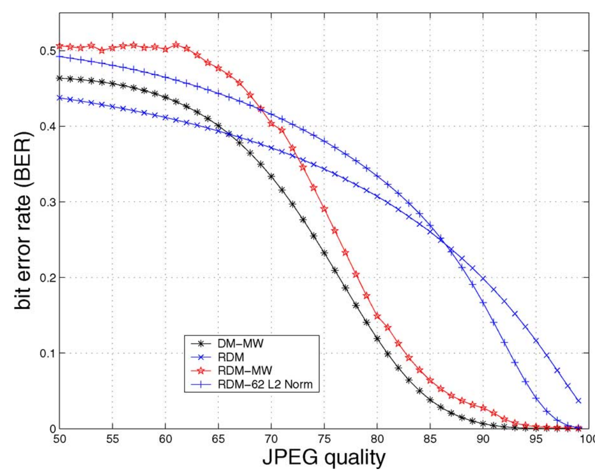

Fig. 22. BER as a function of JPEG quality for a fixed Watson distance of 55.

Fig. 19 show the BER as a function of constant luminance change. The RDM–MW performs better than RDM or DM–MW, although all three algorithms are more sensitive to constant luminance changes. RDM-62-L2-Norm usually has better performance, which is probably attributable to its window size of 62.

Fig. 20 illustrates the response to additive white Gaussian noise. The RDM–MW, RDM-62-L2-Norm, and DM–MW all perform very similarly. Once again, our implementation of RDM with a window size of 1 performs significantly worse.

The sensitivity to JPEG compression is investigated in Figs. 21 and 22. For a fixed DWR of 35 dB, we see that both RDM–MW and DM–MW perform worse that RDM algorithms with no perceptual model. However, if we fix the perceptual distortion rather than the DWR, then the perceptual-based algorithms have superior performance. The RDM–MW has slightly worse performance than DM–MW, which is probably due to the fact that the perceptual model for RDM–MW is worse.

X. Conclusion

We have proposed several modifications to dither modulated QIM. The first method, DM-W, uses Watson's perceptual model to adaptively change the quantization step size in order to improve fidelity. Experimental results confirm that for the same DWR, the Watson distance is reduced by more than 80%. This improvement is achieved while simultaneously improving the robustness in high noise regimes.

Next, we modified Watson's perceptual model so that the adaptive QIM scheme (DM–MW) is theoretically invariant to valumetric scaling. Experimental results demonstrate that using soft-decision decoding, the BER does not exceed 7% over a scale range of 0.5 to 1.5. While there is a small degradation in fidelity compared with DM–W, the perceptual distortion introduced by this method is much lower than the original QIM, the adaptive method of Oostveen *et al.*, and standard rational dither modulation.

This algorithm—DM–MW—implicitly assumes that the slacks calculated at the embedder and detector are the same, despite the modifications to the DCT coefficients introduced by the embedder. This is a source of error that is eliminated by using rational dither modulation with the modified Watson measure (RDM–MW). The adaptive step size for the current block is now based on perceptual estimates from the previously watermarked, neighboring block. This guarantees that the slacks calculated at the embedder and detector are identical (prior to any distortions between embedding and detecting). Of course, using the perceptual slacks calculated from the previously watermarked neighborhood to affect the step size in the current block is likely to introduce some distortions, and a further small degradation in fidelity, compared with both DM–W and DM–MW, is observed. However, this degradation is small and the use of RDM–MW significantly reduces the BER.

Experimental results also demonstrated that the performance of DM–W, DM–MW, and RDM–MW degrade more smoothly with the addition of white Gaussian noise. This provides superior performance in high noise regimes.

We also tested the performance of these algorithms to the addition/subtraction of a constant luminance value. Experimental results comparing the performance with that of Oostveen *et al.* indicate all three algorithms are significantly more robust. However, standard DM in the DCT domain (DM) is the most robust.

Sensitivity to JPEG compression was also investigated. We observed that all of the algorithms considered here are very sensitive to JPEG compression. If we maintain a fixed DWR, we generally observe that the perceptually-based algorithms perform worse that those that do not use a perceptual model. However, if we fix the perceptual distortion rather than DWR, then the perceptually-based algorithms provide significant improvements in performance.

The algorithms we have described are block based. As such, they are robust to spatially varying amplitude changes if such changes occur on a block-by-block basis. However, highly nonlinear amplitude changes, such as gamma correction, remain problematic. This remains an area of future work.

Our experimental investigations also reveal that a significant source of residual errors is due to rounding which occurs after computing the inverse DCT and converting to integer pixel values. These rounding errors mean that the step size calculated at the embedder and detector are not identical. Rounding errors are, of course, a form of quantization, and most QIM, DM, and DC–DM methods are sensitive to requantization which can also occur as a result of analog-to-digital conversion and JPEG compression. The algorithms proposed here have the same susceptibility. Future work will investigate spread-transform dither modulation as an avenue for further improvement.

Acknowledgment

The authors would like to thank J. C. Oostveen, T. Kalker, and M. Staring for providing the code to their algorithm. They would also like to thank G. Doerr for suggesting randomization of the message code.

References

[1] I. J. Cox, M. L. Miller, and A. McKellips, "Watermarking as communications with side information," *Proc. IEEE*, vol. 87, no. 7, pp. 1127–1141, Jul. 1999.

[2] M. Costa, "Writing on dirty paper," *IEEE Trans. Inf. Theory*, vol. IT-29, no. 3, pp. 439–441, May 1983.

[3] B. Chen and G. Wornell, "An information-theoretic approach to the design of robust digital watermarking systems," presented at the Int. Conf. Acoustics, Speech and Signal Processing, Phoenix, AZ, Mar. 1999.

[4] A. B. Watson, "DCT quantization matrices optimized for individual images," *Human Vision, Visual Processing, and Digital Display IV*, vol. SPIE-1913, pp. 202–216, 1993.

[5] J. Eggers, R. Bauml, and B. Girod, "Estimation of amplitude modifications before SCS watermark detection," in *Proc. SPIE: Security, Steganography, and Watermarking of Multimedia Contents*, San Jose, CA, Jan. 2002, pp. 387–398.

[6] K. Lee, D. S. Kim, T. Kim, and K. A. Moon, "Em estimation of scale factor for quantization-based audio watermarking," presented at the Digital Watermarking, 2nd Int. Workshop, Seoul, Korea, Oct. 2003.

[7] J. Oostveen, A. Kalker, and M. Staring, "Adaptive quantization watermarking," in *Proc. SPIE:Security, Steganography, and Watermarking of Multimedia Contents VI*, San Jose, CA, Jan. 2004, vol. 5306, pp. 37–39.

[8] P. Bas, "A quantization watermarking technique robust to linear and non-linear valumetric distortions using a fractal set of floating quantizers," presented at the 7th Information Hiding Workshop, Barcelona, Spain, Jun. 2005.

[9] R. L. Lagendijk and I. D. Shterev, "Estimation of attacker's scale and noise variance for QIM-DC watermark embedding," presented at the IEEE Int. Conf. Image Processing, Singapore, Oct. 2004.

[10] I. D. Shterev, I. L. Lagendijk, and R. Heusdens, "Statistical amplitude scale estimation for quantization-based watermarking," in *Proc. SPIE:Security, Steganography, and Watermarking of Multimedia Contents VI*, San Jose, CA, Jan. 2004, vol. 5306, pp. 796–804.

[11] Q. Li and I. J. Cox, "Using perceptual models to improve fidelity and provide invariance to valumetric scaling for quantization index modulation watermarking," presented at the IEEE Int. Conf. Acoustics, Speech and Signal Processing, Philadelphia, PA, Mar. 2005.

[12] F. Perez-Gonzalez, M. Barni, A. Abrardo, and C. Mosquera, "Rational dither modulation: A novel data-hiding method robust to value-metric scaling attacks," presented at the IEEE Int. Workshop on Multimedia Signal Processing, Siena, Italy, Sep. 2004.

[13] B. Chen and G. W. Wornell, "Quantization index modulation methods for digital watermarking and information embedding of multimedia," *J. VLSI Signal Process. Syst. Signal, Image, Video Technol., Special Issue Multimedia Signal Process.*, pp. 7–33, Feb. 2001.

[14] R. Zamir and M. Feder, "On lattice quantization noise," *IEEE Trans. Inf. Theory*, vol. 42, no. 4, pp. 1152–1159, Jul. 1996.

[15] N. S. Jayant and P. Noll, *Digital Coding of Waveforms: Principles and Applications to Speech and Video*. Englewood Cliffs, NJ: Prentice-Hall, 1984.

[16] I. J. Cox, M. L. Miller, and J. A. Bloom, *Digital Watermarking*. San Mateo, CA: Morgan Kaufmann, 2001.

[17] A. B. Watson, "Visually optimal DCT quantization matrices for individual images," in *Proc. IEEE Data Comp. Conf.*, 1993, pp. 178–187.

[18] Q. Li and I. J. Cox, "Rational dither modulation watermarking using a perceptual model," in *Proc. IEEE Int. Workshop Multimedia Signal Processing*, Shanghai, China, Oct. 2005.

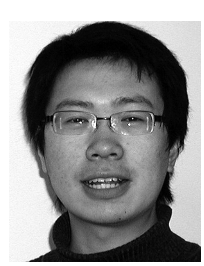

Qiao Li (S'05) received the B.Sc. and M.Sc. degrees in electronic science and technology from Harbin Institute of Technology, Harbin, China, in 2000 and 2002, respectively, and is currently pursing the Ph.D. degree in the Department of Electrical and Electronic Engineering at University College London, London, U.K.

He was a VLSI Design Engineer with NTT Electronics, Tokyo, Japan, from 2001 to 2002. His research interests include information security, multimedia and communications, digital signal processing, and VLSI design. He received the U.K. Overseas Research Students (ORS) Awards Scheme and UCL Graduate School Fellowship.

Ingemar J. Cox (F'06) received the B.Sc. degree from the University College London, London, U.K. in 1980 and the Ph.D. degree from Oxford University, Oxford, U.K., in 1983.

Currently, he is Professor and BT Chair of Telecommunications in the Departments of Electronic Engineering and Computer Science at University College London and Director of UCL's Adastral Park Postgraduate Campus. He was a member of the Technical Staff at AT&T Bell Labs, Murray Hill, NJ, from 1984 to 1989, where his research interests were focused on mobile robots. In 1989, he joined the NEC Research Institute, Princeton, NJ, as a Senior Research Scientist in the computer science division. At NEC, his research shifted to problems in computer vision and he was responsible for creating the computer vision group at NECI. He has worked on problems involving stereo and motion correspondence and multimedia issues of image database retrieval and watermarking. From 1997 to 1999, he was Chief Technical Officer of Signafy, Inc., Princeton, NJ, a subsidiary of NEC responsible for the commercialization of watermarking. Between 1996 and 1999, he led the design of NEC's watermarking proposal for DVD video disks and later colloborated with IBM in developing the technology behind the joint "Galaxy" proposal supported by Hitachi, IBM, NEC, Pioneer, and Sony. In 1999, he returned to the NEC Research Institute as a Research Fellow. He is coauthor of a book *Digital Watermarking* and the co-editor of two books *Autonomous Robots Vehicles* and *Partitioning Data Sets: With Applications to Psychology, Computer Vision and Target Tracking*.

Dr. Cox is a Fellow of the Institution of Electrical Engineers (IEE) and a Fellow of the Royal Society for Arts and Manufacturers. He is a Holder of a Royal Society Wolfson Fellowship. He is Co-Editor-in-Chief of the *IEE Proceedings on Information Security* and an Associate Editor of the IEEE TRANSACTIONS ON INFORMATION FORENSICS AND SECURITY. In 1999, he was awarded the IEEE Signal Processing Society Best Paper Award (Image and Multidimensional Signal Processing Area) for a paper he coauthored on watermarking.