Toward Robust Logo Watermarking Using Multiresolution Image Fusion Principles

Deepa Kundur, *Senior Member, IEEE,* and Dimitrios Hatzinakos, *Senior Member, IEEE*

***Abstract*—This paper presents a novel robust watermarking approach called FuseMark based on the principles of image fusion for copy protection or robust tagging applications. We consider the problem of logo watermarking in still images and employ multiresolution data fusion principles for watermark embedding and extraction. A human visual system model based on contrast sensitivity is incorporated to hide a higher energy hidden logo in salient image components. Watermark extraction involves both characterization of attacks and logo estimation using a rake-like receiver. Statistical analysis demonstrates how our extraction approach can be used for watermark detection applications to decrease the problem of false negative detection without increasing the false positive detection rate. Simulation results verify theoretical observations and demonstrate the practical performance of FuseMark.**

***Index Terms*—Data fusion, digital watermarking, image tagging, logo watermarking, multimedia security.**

I. Introduction

In THIS PAPER, we focus on the signal processing aspects of the digital watermarking problem for the purpose of robust grayscale logo embedding. We determine strategies to improve watermark robustness through the incorporation of image fusion tool-sets. We propose a multiresolution image watermarking technique called *FuseMark* that addresses the problem of invisible logo embedding in still images when the original *host* signal is available for watermark extraction. Although the host accessibility assumption is somewhat impractical for many data hiding situations, we consider applications in which an automated search is employed by a party with access to the host. These types of watermarking techniques, termed *nonblind*, have shown much higher robustness to intentional attacks than their *blind* counterparts. Early work in the area of robust nonblind watermarking involved spread spectrum technologies applied in various domains and often employing perceptual models [1]–[7].

The success of watermarking for intellectual property management depends, in part, on how easily it is adopted within common policy, and on how effectively infringement cases are addressed [8]. It has been suggested that *visually meaningful* watermark images may improve the trustworthiness of signal tracing, identification or security based on data embedding in the eyes of nontechnical arbitrators [9]. The presentation of a recognizable mark is much more convincing than a numerical value and allows the opportunity to exploit the human visual system's ability to recognize a pattern. As discussed in [10], in the same way channel coding improves robustness against transmission errors, the visual perception process naturally filters out random noise to better recognize a meaningful pattern. This has motivated us to research embedding logos into images. Since a logo is considered itself to be a small image, we make use of image fusion architectures for watermarking instead of the well-established SS strategies discussed previously. SS techniques are more useful for detecting the presence of a given watermark, than for extracting an arbitrary logo.

Grayscale logo watermarking has been applied primarily to visible watermarking. In [11], Braudway *et al.* use analytic human perceptual models to embed, by varying pixel brightness, a reasonably unobtrusive visible logo in still color images; the process of adjusting the watermark intensity for a given image is automated in [12] by using a texture-based HVS metric. In [13], Meng and Chang borrow ideas from [11] to extend visible watermarking for DCT-based video streams. Similarly, Kankanhalli *et al.* in [14] work in the frequency domain to classify segmented 8×8 DCT image blocks into one of eight regions exhibiting different masking characteristics. The visible logo is embedded in the luminance color component and the associated watermark scaling factors for each DCT block are based on the associated masking classification and empirical formulas employing luminance sensitivity.

Invisible binary logo marking has also been investigated by a number of researchers. In [15], Voyatzis and Pitas use a chaotic system to securely scramble a binary watermark into a pseudo-random bit sequence before embedding the information by modifying the intensity values of selected host image pixels. Similarly, in [16], Lin converts a binary logo into a mark sequence using a unique serial number and embeds by changing pixel block intensities in a novel way. Su and Kuo [17] also design a scheme to detect binary logos, but in cartoon and map images; the mark embedding process involves modifying the bit planes of DWT coefficient subbands. In [18], the authors present a nonblind method to embed a gray-level image by first converting the watermark into bit planes. The individual bit planes are then inserted as a set of binary logos. More recently, Zeng *et al.* [10], consider resolving rightful ownership through embedding a binary image and present

Manuscript received November 21, 2000; revised June 4, 2002. This work was supported by the Natural Sciences and Engineering Research Council (NSERC) and the Communications and Information Technology Ontario (CITO). The associate editor coordinating the review of this paper and approving it for publication was Dr. Minerva Yeung.

D. Kundur was with the Edward S. Rogers Sr. Department of Electrical and Computer Engineering, University of Toronto, Toronto, ON M5S 3G4 Canada. She is now with the Department of Electrical Engineering, Texas A&M University, College Station, TX 77843-3128 USA (e-mail: deepa@ee.tamu.edu).

D. Hatzinakos is with the Edward S. Rogers Sr. Department of Electrical and Computer Engineering, University of Toronto, Toronto, ON M5S 3G4 Canada (e-mail: dimitris@comm.toronto.edu).

Digital Object Identifier 10.1109/TMM.2003.819747

ways to visualize the extracted mark. Adjacent logo bits are then inserted into adjacent host image blocks. For improved robustness in high degree attack situations, adjacent image blocks are incorporated into the extraction procedure of a given logo bit to produce a lower resolution, but visually appealing extracted watermark result.

In [19], Knox proposes the concept of a reversible image for data hiding, in which two 8-bit grayscale images are fused into a signal 8-bit grayscale image. Essentially, the four most significant bits (MSBs) of one image are placed in the corresponding four MSBs of the new image. The four MSBs of the second image are hidden in the least significant bits (LSBs) of the new image. The combined result is called reversible because swapping the order of the bits, so that the LSBs become the MSBs and vice versa makes the second image viewable while masking the first.

By treating the logos as bit sequences, the proposed techniques cannot conveniently account for the perceptual characteristics of both the host and the watermark for effective data hiding. Considering arbitrary grayscale intensity logos facilitates the embedding of arbitrary commercial logos and increases the quality of and overall number of possible logos identifiable by human observers. This motivates us to use image fusion principles for *invisible robust grayscale* logo watermarking.

The process of embedding and extracting the picture logo is analogous to image fusion in a number of ways.

- Both processes, to be effective, must make use of sophisticated human visual models. For watermarking, the data is hidden from perception; for fusion, the merged information must be displayed so that the *salient* components of each image are apparent upon viewing.
- Watermarking and fusion are both information compression problems in which two or more separate pieces of data are combined into one result.
- Robust logo extraction involves identifying and merging noisy versions of the embedded watermark from various components of the host image which can be viewed as an image fusion problem.

A. Contributions and Scope

The objectives of this paper are to

1) present a watermarking scheme designed with the use of image fusion tool-sets. The technique makes use of the discrete wavelet transform (DWT), a novel perceptual weighting model to embed the watermark, and repetition, for combined robustness;
2) develop a method to estimate the logo watermark from a possibly attacked watermarked image to maximize signal-to-noise ratio (SNR) of the extracted logo. We employ a minimum variance fusion approach;
3) analyze the performance of the proposed fusion-based algorithm for watermark detection applications. We demonstrate how the proposed approach reduces the probability of false negative detection without increasing the false positive detection rate.

In Section II, we discuss the proposed technique. Section III looks at performance advantages for watermark detection applications. Simulation results are presented in Section IV followed by a discussion and conclusions in Sections V and VI, respectively.

II. Multiresolution Fusion-Based Watermarking

A. Why Multiresolution Fusion?

Image fusion is a process that produces a single image from a set of input images. The fused image should contain more "complete" information (as defined by a given application), than any individual input [20]. Since this is a sensor-compressed information problem, it follows that wavelets, classically useful for HVS processing, data compression and reconstruction, are also helpful for such merging [21]–[23]. Some multiresolution wavelet fusion methods make use of HVS models to determine the perceptually most significant information from each image to retain in the composite [21]. It is then expected that such rules can be applied to judiciously select the components of the host image in which to embed the watermark.

B. Fusemark

A multiresolution data fusion approach is employed in which the image and the watermark are transformed into the discrete wavelet domain. The resulting coefficients are then *fused* according to a series of combination rules that take into account the HVS characteristics. The watermark is restricted to be much smaller in dimension than the host signal, so that it can be repeatedly fused throughout the signal. Because the fusion process takes place at different resolution levels and at various localized spatial regions, the watermark is spread throughout the different frequencies of the host, but is spatially localized at higher resolutions of the host image. This provides robustness against varying forms of signal distortions including filtering, subsampling and cropping.

A preliminary version of this fusion-based watermarking algorithm was first presented in [24]. In this paper, we extend our scheme by incorporating a minimum variance fusion/rake-like receiver for watermark estimation and provide statistical performance analysis.

1) The Three Stage Method: Throughout our discussion, we use $f(m,n)$ to denote the host image and $w(m,n)$ the watermark, assumed to be a two-dimensional array of real elements. We assume that the size of the watermark is smaller than the host by a factor of 2^{n_x} and 2^{n_y} in the x and y dimensions, respectively, where $n_x, n_y \in \mathbb{Z}^+$ (where $\mathbb{Z}^+$ represents the set of all integers greater than zero), and that the dimensions of the watermark are $2N_{wx} \times 2N_{wy}$. *These conditions facilitate our description and analysis of the technique, however the dimensions are not practically restricted to these forms.*

The general technique is summarized in Fig. 1. The Sections II-B2 and II-B3 provide a more detailed and analytic description.

Stage 1: We perform the Lth level DWT of the host image to produce a sequence of $3L$ *detail* images at each of the L resolution levels, and a gross approximation of the image at the coarsest resolution level. The value of L is user-specified and is set to a positive integer less than or equal to $\min(n_x, n_y) - 1$. We denote the oth frequency orientation at the lth resolution level

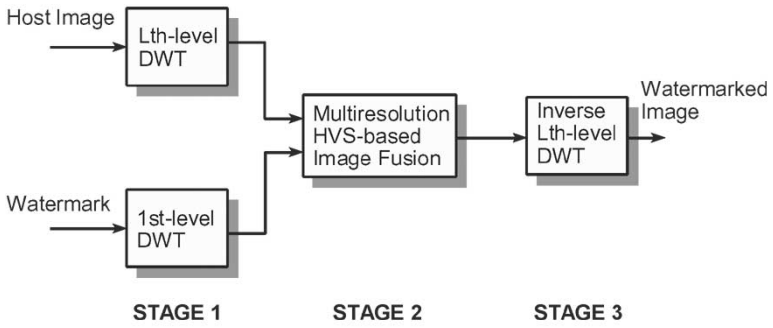

Fig. 1. Proposed fusion-based watermark embedding method.

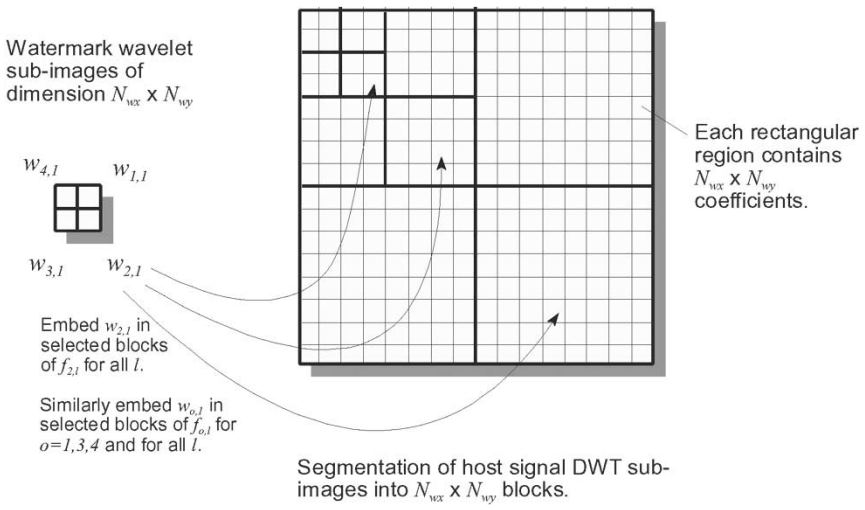

Fig. 2. Segmentation of the host image wavelet coefficients into $N_{wx} \times N_{wy}$ blocks for fusion watermarking. The salience of each block is computed and if it is above a threshold specified by B, then the corresponding $N_{wx} \times N_{wy}$ watermark wavelet coefficient is embedded. As suggested by the diagram, the watermark is more widely spread spatially when embedded at a lower (coarser) resolution.

of the image f by $f_{o,l}(m,n)$ where $o \in \{1,2,3\}$ represents the frequency orientation corresponding to the horizontal, diagonal and vertical image details, $l \in \{1,2,\ldots,L\}$ the resolution level and (m,n) the particular spatial location index at the resolution l. The gross approximation is represented by $f_{4,L}(m,n)$ where the subscript "4" is used instead of o to denote the gross image approximation at resolution L.

Similarly, the first-level DWT of the watermark w is performed to produce $N_{wx} \times N_{wy}$ dimensional detail and approximation subimages $w_{o,1}(m,n)$, $o \in \{1,2,3,4\}$.

Stage 2: The subimages, $f_{o,l}(m,n)$ (for $o = 1,2,3$ and $l = 1,2,\ldots,L$) and $f_{4,L}(m,n)$ are segmented into adjacent nonoverlapping $N_{wx} \times N_{wy}$ blocks. Fig. 2 demonstrates the procedure. We denote these rectangular blocks by $f^i_{o,l}(m,n)$, $i = 1,\ldots,2^{n_x+n_y-2l}$, where it can be shown that $2^{n_x+n_y-2l}$ is the total number of $N_{wx} \times N_{wy}$ blocks at each frequency orientation o and resolution l.

The *salience*, $\mathcal{S}$, which is a numerical measure of perceptual importance, of each of these localized segments is computed using a model of the *contrast sensitivity* of the HVS. Contrast sensitivity is defined as the reciprocal of the contrast necessary for a given spatial frequency to be perceived. Experimental tests have resulted in the well-known model given by Dooley [25]. The resulting contrast sensitivity for a particular pair of spatial frequencies is given by:

$$C(\omega_1,\omega_2) = 5.05e^{-0.178(\omega_1+\omega_2)}(e^{0.1(\omega_1+\omega_2)} - 1), \qquad (1)$$

where $C(\omega_1,\omega_2)$ is the contrast sensitivity matrix and ω_1 and ω_2 are the spatial frequencies given in units of cycles per visual angle (in degrees). A conversion from cycles per visual angle to radians per pixel must be made prior to the use of C in our algorithm. For the simulations in this chapter, we apply a conversion assuming a 256×256 host image and a viewing distance of six times the image size [21].

Our definition of salience, first proposed and tested in [21] for perceptually-based image fusion, provided a numerical value of visual importance and, is given by

$$\mathcal{S}(f^i_{o,l}(m,n)) = \sum_{\forall\ (\omega_1,\omega_2)} C(\omega_1,\omega_2)|F^i_{o,l}(\omega_1,\omega_2)|^2 \qquad (2)$$

where C is the contrast sensitivity matrix, and $F^i_{o,l}(\omega_1,\omega_2)$ is the normalized discrete Fourier transform of the image component $f^i_{o,l}(m,n)$; $F^i_{o,l}(\omega_1,\omega_2)$ is normalized such that it has unit energy (i.e., $\|F^i_{o,l}(\omega_1,\omega_2)\|^2 = 1$).

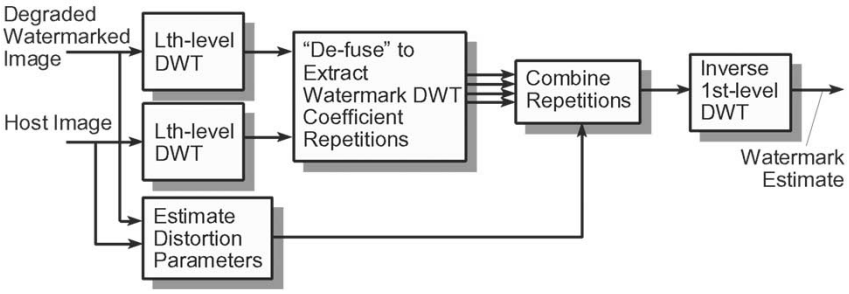

Fig. 3. Proposed fusion-based watermark extraction method.

The watermark is embedded only in B percent[1] of the most salient detail coefficient image blocks at each resolution level and orientation using the following equation:

$$f_{o,l}^{w,i}(m,n) = f_{o,l}^{i}(m,n) + s_{o,l}^{i} w_{o,1}(m,n) \tag{3}$$

where $f_{o,l}^{w,i}(m,n)$ is the watermarked DWT coefficient. For the remaining blocks, we set

$$f_{o,l}^{w,i}(m,n) = f_{o,l}^{i}(m,n). \tag{4}$$

The parameters $s_{o,l}^{i}$ are positive real numbers that determine a tradeoff between the visibility of the watermark and its robustness to signal distortion at each of the resolution levels. The following rule of thumb has been determined to set these parameter values:

$$s_{o,l}^{i} = \overline{s} \sum_{\forall\,(m,n)} |f_{o,l}^{i}(m,n)| \sqrt{\frac{\mathcal{S}(f_{o,l}^{i}(m,n))}{\max_{\text{over all } j} \mathcal{S}(f_{o,l}^{j}(m,n))}}, \tag{5}$$

where the absolute scale $\overline{s}$ ranges between 0.2 and 1 for our test photographs. The fraction within the square root of (5) is a relative measure that gives greater weight judiciously to the embedded watermark in more salient host image regions.

A similar merging procedure is used to embed the watermark approximation coefficient $w_{4,1}(m,n)$ into the host image approximation blocks $f_{4,L}^{i}(m,n)$. The watermark is embedded in all blocks and $\tilde{s}$ (which is used instead of $\overline{s}$ to distinguish the two parameters) is set between 0.02 and 0.2 to ensure imperceptibility.

Stage 3: The corresponding Lth-level inverse DWT of the fused image components $f_{o,l}^{w}(m,n)$ is computed to form the watermarked image.

The parameter L determines the maximum resolution at which the watermark is embedded. In general, the larger the value of L, the more localized the watermark is in the lower frequencies of the host image that makes it naturally robust to distortions which affect only image details.

2) Weighted Watermark Extraction: A summary of the watermark extraction process is provided in Fig. 3. To extract the watermark, the DWT is applied to the potentially degraded watermarked image. We denote the resulting coefficients $\hat{f}_{o,l}^{w,i}(m,n)$. Each repetition of the watermark DWT coefficients is extracted through subtraction of the host to produce the estimates $\hat{w}_{o,l}^{i}(m,n)$ as follows:

$$\hat{w}_{o,l}^{i}(m,n) = \frac{\hat{f}_{o,l}^{w,i}(m,n) - f_{o,l}^{i}(m,n)}{s_{o,l}^{i}}. \tag{6}$$

We choose to model the attacks on the watermarked image blocks as locally stationary AWGN. This provides opportunity to adapt to the particular image distortions to reliably recompose the watermark. In addition, the model involves only one simple locally invariant parameter to reflect the noise power at a given image location. The limitation is that our degradation model excludes geometric attacks such as rotation; effects of nongeometric attacks such as compression, filtering and noise, however, fall under the umbrella of our model. Specifically, each corresponding DWT coefficient image block $\hat{f}_{o,l}^{w,i}(m,n)$ of the attacked watermarked image for $l = 1, 2, \ldots, L$, and $o = 1, 2, 3$, is assumed to undergo degradation of the form

$$\hat{f}_{o,l}^{w,i}(m,n) = f_{o,l}^{w,i}(m,n) + v_{o,l}^{i}(m,n) \tag{7}$$

in which $v_{o,l}^{i}(m,n)$ is zero mean AWGN with known variance $\sigma^2_{v_{o,l}^{i}}$. The local characterization of the distortions allows more optimal combining of the watermark repetitions.

Similarly, we assume for the gross approximation coefficients $\hat{f}_{4,L}^{w,i}(m,n)$ that

$$\hat{f}_{4,L}^{w,i}(m,n) = f_{4,L}^{w,i}(m,n) + v_{4,L}^{i}(m,n) \tag{8}$$

where $v_{4,L}^{i}(m,n)$ is zero mean AWGN with known variance $\sigma^2_{v_{4,L}}$, which is independent of the block position i. Distortions that attempt to maintain the perceptual quality of the watermarked image generally have a less diverse effect on the amplitude of these coefficients, so we neglect any block variation in the noise power. Simulation results also demonstrate superior behavior when ignoring these small variations.

Equation (7) suggests that $\hat{w}_{o,l}^{i}(m,n)$ can be represented by

$$\hat{w}_{o,l}^{i}(m,n) = w_{o,1}(m,n) + \frac{v_{o,l}^{i}(m,n)}{s_{o,l}^{i}} \tag{9}$$

$$= w_{o,1}(m,n) + u_{o,l}^{i}(m,n) \tag{10}$$

where $u_{o,l}^{i}(m,n)$ is zero mean AWGN with variance $\sigma^2_{v_{o,l}^{i}}/(s_{o,l}^{i})^2$. Thus, our extracted repetitions are noisy versions of the originals, where the noise power varies for each repetition. Taking a data fusion perspective, we can combine the repetitions by applying minimum variance fusion based

[1]The larger B, the greater the visibility of the watermark. In our tests $25 \leq B \leq 75$ allows for an appropriate tradeoff between perceptibility and robustness.

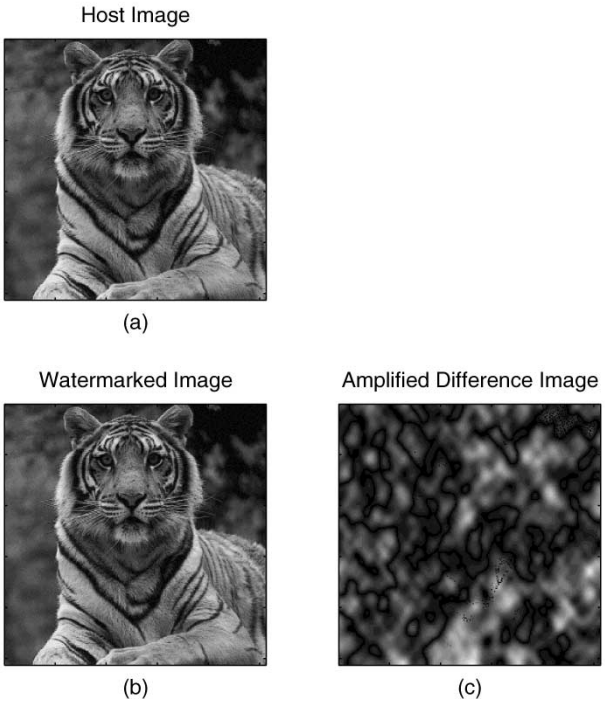

Fig. 4. Results for the DCT method: (a) host image, (b) watermarked image, and (c) amplified absolute difference between the images of (a) and (b).

on the principles of Kalman filtering [20]. The estimate of the DWT watermark coefficients $\hat{w}^i_{o,l}(m,n)$ is produced as

$$\hat{w}_{o,1}(m,n) = \sum_{l=1}^{L} \sum_{i \in \mathbb{S}_{o,l}} \alpha^i_{o,l} \hat{w}^i_{o,l}(m,n) \tag{11}$$

for $o = 1,2,3$, where $\mathbb{S}_{o,l}$ is the set of all block indices containing an embedded mark, and $\alpha^i_{o,l}$ is given by

$$\alpha^i_{o,l} = \frac{\left(\frac{s^i_{o,l}}{\sigma_{v^i_{o,l}}}\right)^2}{\sum_{l=1}^{L} \sum_{j \in \mathbb{S}_{o,l}} \left(\frac{s^j_{o,l}}{\sigma_{v^j_{o,l}}}\right)^2}. \tag{12}$$

The details of the derivation of $\alpha^i_{o,l}$ are provided in [26]. Similarly,

$$\hat{w}_{4,1}(m,n) = \sum_{i=1}^{2^{n_x+n_y-2L}} \frac{1}{2^{n_x+n_y-2L}} \hat{w}^i_{4,L}(m,n). \tag{13}$$

Equations (11) and (13) resemble the rake receiver structure with maximal ratio combining used in multipath communication environments [27], [28]. This SNR maximization procedure helps in the estimation of both logo-type watermarks and statistical watermarks used for detection. Finally, the inverse DWT is computed on the coefficient estimates to form the overall extracted watermark $\hat{w}$.

3) On the Estimation of Distortion Parameters: The calculation of the optimal weights $\alpha^i_{o,l}$ requires knowledge of noise variances $\{\sigma^2_{v^i_{o,l}}\}$ of each block containing the watermark. We estimate the noise variance parameter of block i from the closest adjacent block not containing a watermark. Specifically, the parameters are computed as follows:

$$\hat{\sigma}^2_{v^i_{o,l}} = \frac{1}{N_{wx} N_{wy}} \sum_{\forall\,(m,n)} |\hat{f}^{w,i^*}_{o,l}(m,n) - f^{i^*}_{o,l}(m,n)|^2, \tag{14}$$

where $i \in \mathbb{S}_{o,l}$ and i^* is the spatially closest block to i at the same resolution in which no watermark is embedded in the region (i.e., $i^* \notin \mathbb{S}_{o,l}$).

The reader should note that we do not make use of the watermark w during the calculation of the weights. This quantity may not be known or its use in computing $\alpha^i_{o,l}$ may result in a biased watermark estimate; the details are discussed in Section III.

III. Watermark Detection

Although the algorithm has been designed with logos in mind, elements such as weighted watermark extraction have advantages for statistical watermark detection as well. For watermark detection a randomly generated noiselike sequence $w(i)$ is embedded. To test the existence of a particular watermark, the corresponding extracted sequence is correlated with the known embedded mark as follows

$$\rho(w,\hat{w}) = \frac{\sum_{i=1}^{N_w} w(i)\hat{w}(i)}{\sqrt{\sum_{i=1}^{N_w} w^2(i)} \sqrt{\sum_{i=1}^{N_w} \hat{w}^2(i)}}, \tag{15}$$

where w and $\hat{w}$ are the known and extracted watermarks, respectively. If and only if $\rho(w,\hat{w}) \geq \mathcal{T}$, where $\mathcal{T} > 0$ is a prespecified threshold, the watermark is positively detected.

Due to the varying characteristics of the distortion in an image, many techniques such as [7] extract the watermark separately from the different image resolution and frequency orientations. Each extracted watermark segment is separately correlated with the embedded and the maximum value of the correlation is used for watermark detection. We call this *maximum correlation detection*. In this way, highly degraded portions of the watermark are not used for detection. Although this approach increases robustness, it also increases the probability of false watermark detection as we show in Section III-A. For commercial watermarking applications such as DVD copy protection, an increased threat of false detection is of serious concern and may not be worth the performance improvement [29].

We next demonstrate how the use of our adaptive weighted watermark extraction method improves performance without increasing the probability of false positive detection.

A. Probabilities of False Detection

For simplicity we assume that M different repetitions of a randomly generated binary watermark signal are embedded within the host and that there is no host signal interference since the algorithm is nonblind. We summarize the results derived in [26] for the following.

1) *Elementary detection* in which an arbitrary repetition of the watermark is extracted and correlated with the known mark for detection.

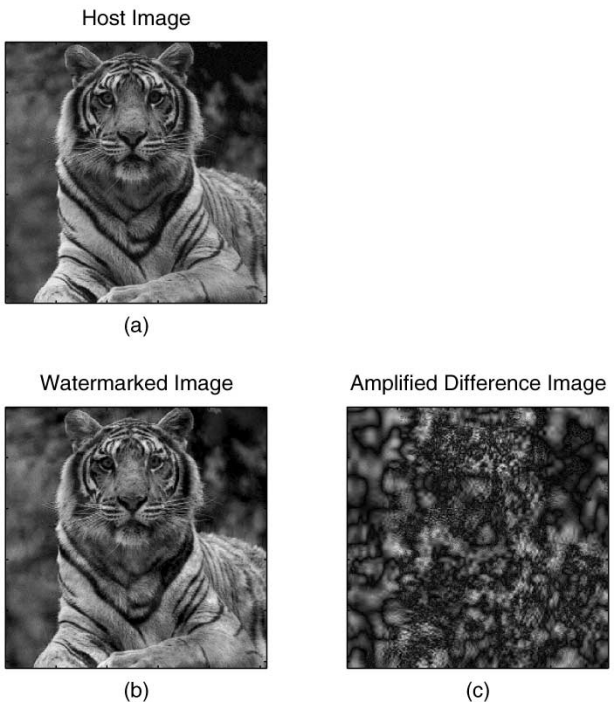

Fig. 5. Results for the proposed multiresolution fusion-based method: (a) host image, (b) watermarked image, and (c) amplified absolute difference between the images of (a) and (b).

2) *Maximum correlation* detection, in which each repetition is separately extracted and correlated. The maximum correlation value is used for detection.
3) *Optimal weight detection*, which we described in Section II-B2, in which the watermark repetitions are individually weighted to improve the overall SNR of the extracted watermark. Then, correlation is performed.

1) Elementary Detection: The probability of false negative, P_{fn}, and the probability of false positive, P_{fp}, watermark detections are calculated to be [26]

$$P_{fn} \approx \frac{1}{2}\text{erfc}\left(\sqrt{\frac{N_w SNR_{\hat{w}}}{2}}(1-\mathcal{T})\right) \tag{16}$$

and

$$P_{fp} \approx \frac{1}{2}\text{erfc}\left(\sqrt{\frac{N_w}{2}}\mathcal{T}\right), \tag{17}$$

respectively, where erfc is the complementary error function, N_w is the length of the watermark, $\mathcal{T}$ is the prespecified threshold for detection and $SNR_{\hat{w}}$ is the SNR of the extracted watermark. Our analysis assumes that the extracted watermark exhibits zero mean AWGN with finite variance.

2) Maximum Correlation Detection: The probability of false negative watermark detection can be lowered by employing maximum correlation detection. As derived by the first author of this paper in [26], P_{fn} and P_{fp} are

$$P_{fn} \approx \prod_{k=1}^{M}\frac{1}{2}\text{erfc}\left(\sqrt{\frac{N_w SNR_{\hat{w},k}}{2}}(1-\mathcal{T})\right) \tag{18}$$

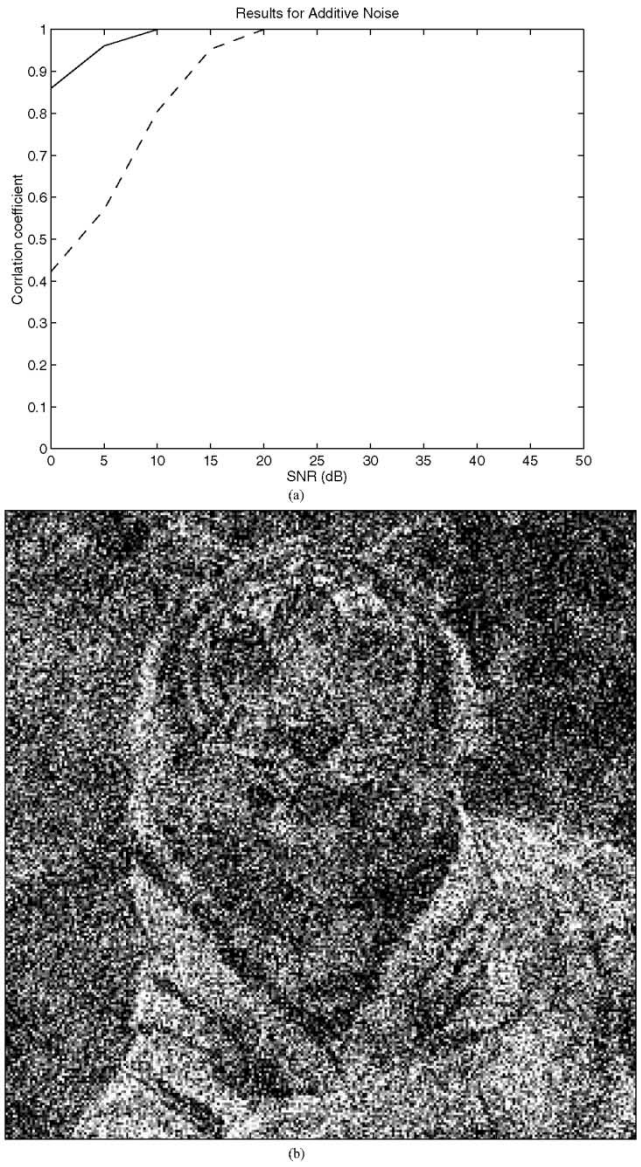

Fig. 6. Results for additive white gaussian noise degradation: (a) correlation coefficient versus SNR (the dashed and solid lines correspond to the DCT and the proposed methods, respectively) and (b) degraded watermarked image at 0 dB SNR.

and

$$P_{fp} \approx \frac{M}{2}\text{erfc}\left(\sqrt{\frac{N_w}{2}}\mathcal{T}\right) \tag{19}$$

respectively, where M is the number of repetitions of the watermark signal separately extracted and correlated from the different regions of the image and $SNR_{\hat{w},k}$ is the associated SNR of the kth repetition.

3) Optimal Weight Detection: In contrast, using the optimal weighting strategy of Section II-B2, we find that

$$P_{fn} \approx \frac{1}{2}\text{erfc}\left(\sqrt{\frac{N_w SNR^{OW}}{2}}(1-\mathcal{T})\right) \tag{20}$$

and

$$P_{fp} \approx \frac{1}{2}\text{erfc}\left(\sqrt{\frac{N_w}{2}}\mathcal{T}\right) \tag{21}$$

where SNR^{OW} is the SNR of the overall weighted watermark estimate. For the specific case in which the noise energy on the watermark is independent of frequency orientation o in which it is embedded, it can be shown that [26]

$$SNR^{OW} = \sum_{k=1}^{M} SNR_{\hat{w},k} \tag{22}$$

where M is the number of repetitions of the watermark signal separately extracted and correlated from the different regions of the image and $SNR_{\hat{w},k}$ is the associated SNR of the kth repetition.

Thus, our postprocessing decreases the probability of false negatives without sacrificing the possibility of a false positive result. The extracted watermark has an effective SNR that is the sum of those of the individual repetitions. This improvement is also highly attractive for logo-type watermarking in which the watermark is merely extracted and displayed rather than passed through a correlator for detection.

IV. Simulation Results

Although the probability of false negative is lower, the probability of false positive increases linearly. We specifically make use of the Daubechies 10-point wavelet [30] that provides a good tradeoff between spatial and frequency localization for all simulations. The well-known DCT method by Cox *et al.* [1] has been implemented for comparison for robustness against watermark detection. Instead of watermarking with an AWGN sequence, we use a randomly generated binary sequence for compatibility with our algorithm. The implementation is as specified in [1] with a "scaling parameter" $a = 0.1$ as suggested in the paper. The DCT method involves adding the watermark to the N_w largest magnitude nondc DCT coefficients of the host image where N_w is the length of the watermark sequence. We embed the ith element of the watermark w according to

$$F_w^{DCT}(\omega_1, \omega_2) = F^{DCT}(\omega_1, \omega_2)\,(1 + aw(i)) \tag{23}$$

where $F^{DCT}(\omega_1, \omega_2)$ and $F_w^{DCT}(\omega_1, \omega_2)$ are the DCT coefficients of the host image f and the watermarked image f_w, respectively, a is the scaling parameter discussed earlier, and $w(i)$ is the ith watermark element. The watermarked image f_w is formed by taking the inverse DCT of the coefficients $F_w^{DCT}(\omega_1, \omega_2)$.

We perform two classes of tests. We first compare the watermark detection capability of our fusion-based method with the DCT method. Randomly generated binary watermarks[2] are embedded within the host signal. The resulting watermarked signal is corrupted using one of many common distortions that we discuss in Section IV-A. The watermark is then extracted and correlated with the embedded watermark sequence to measure detection capability.

[2]The distribution of ones and zeros are assumed to be equiprobable.

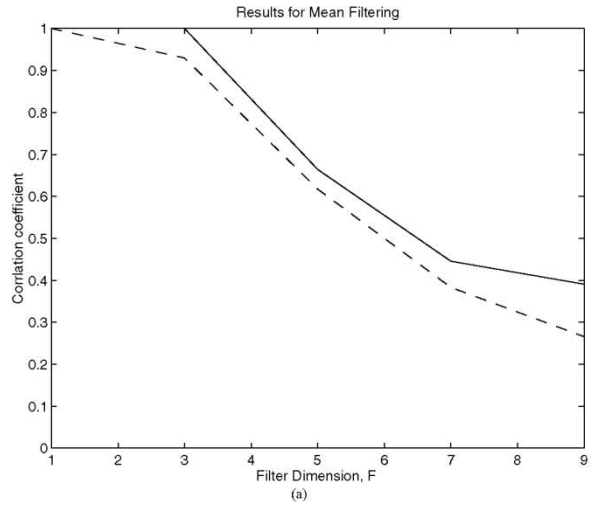

Fig. 7. Results for $F \times F$ mean filtering degradation: (a) correlation coefficient versus dimension of filter F (the dashed and solid lines correspond to the DCT and the proposed methods, respectively) and (b) degraded watermarked image for 9×9 mean filtering.

In the next set of tests, we embed a grayscale logo to demonstrate the performance of the proposed method in embedding and recovering logos when the watermarked image undergoes distortion. We compare the improved performance of using our proposed weighted watermark extraction to our previously proposed technique [24] where simple averaging is performed.

A. Results of Binary Watermark Detection

We perform simulations on the 256×256 host image shown in Fig. 4(a) and a 256 bit randomly generated binary watermark comprised of the elements $\{-1, 1\}$.[3] The corresponding watermarked image using the DCT method with $a = 0.1$ is

[3]We find experimentally that use of $\{-1, 1\}$ instead of $\{0, 1\}$ improves invisibility of the watermark.

displayed in Fig. 4(b). Larger values of a cause a visible change in image contrast. To give an idea of the "shape" of the watermark, we also provide an amplified version of the absolute difference between the watermarked and host images in Fig. 4(c). The watermark is generally smooth and stationary throughout the image.

Similarly, the results for fusion-based watermarking are provided in Fig. 5. The same 256 bit watermark was converted to a 16×16 binary watermark signal for use in the algorithm. The simulations were conducted using the following parameters $\overline{s} = 0.45$, $\tilde{s} = 0.18$, $L = 4$, and $B = 50$. The absolute difference between the watermarked image [shown in Fig. 5(b)] and the host image is displayed in Fig. 5(c). Because of the adaptive and localized nature of our embedding routine the watermark takes on characteristics similar to the host image itself. The use of the DWT and a model of HVS allows the design of an embedded signal that is more naturally masked by the host image itself.

We evaluate and compare the performance of both techniques to various types of image distortions that we discuss below. The correlation coefficient given by (15) between the embedded and extracted watermark is computed to assess robustness. We choose the threshold $\mathcal{T} = 0.4$ that has an associated probability of false positive detection less than 10^{-10} for a 256 bit watermark [26].

The watermarked image was degraded by applying additive white Gaussian noise (AWGN) with varying power. Fig. 6 presents the corresponding correlation coefficients between the embedded and extracted watermarks for different SNRs. The dashed and solid lines correspond to the DCT and the proposed methods, respectively. The proposed technique performs better than the DCT method for lower SNRs. Fig. 6(b) shows the degraded watermarked image (using the proposed method) for an SNR of 0 dB. The watermark correlation coefficient is still high enough to be detected for this degree of distortion.

Figs. 7 and 8 show the results for mean and median filtering, respectively. An $F \times F$ mean (or median) filter was applied on the watermarked image to attempt to destroy the watermark detection capability. The dashed and solid lines correspond to the DCT and the proposed methods, respectively. For larger filter sizes the proposed technique outperforms the DCT method. The watermark detection capability of FuseMark persists even when the image is significantly degraded with 9×9 filters as shown in Figs. 7(b) and 8(b).

The results for JPEG compression are displayed in Fig. 9 for different compression ratios (CR). For low CRs, the DCT and the fusion-based method perform comparably. However, for higher CRs the proposed technique is superior. The watermark is detected using a threshold of 0.4 even for a CR of 35. The resulting compressed image shows visible blocking artifacts as displayed in Fig. 9(b).

The effects of image cropping on watermark detection is shown in Fig. 10. The correlation coefficient as a function of the percentage of the image area remaining is displayed. For watermark extraction, the portion of the watermarked image cropped out was replaced with the host image as performed in [1]. FuseMark has significantly superior performance because of the inherent localization of the watermark at higher resolution levels. Even when only 2.25% of the image area remains, the correlation coefficient for the proposed technique is high.

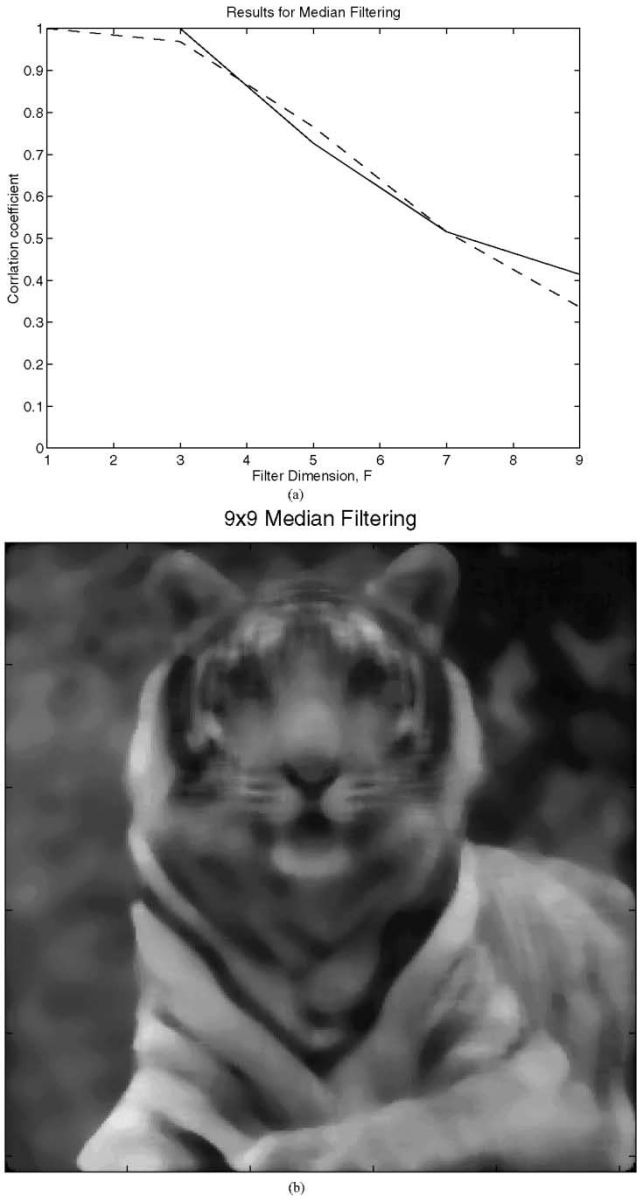

Fig. 8. Results for $F \times F$ median filtering degradation: (a) correlation coefficient versus dimension of filter F (the dashed and solid lines correspond to the DCT and the proposed methods, respectively) and (b) degraded watermarked image for 9×9 median filtering.

Fig. 11 provides the results for resizing of the marked images. The images were scaled down in size by a factor of F using bilinear interpolation and were resized to their original dimensions before watermark extraction. FuseMark performs comparably to the DCT method for smaller values of F. The correlation coefficient for both methods is reasonably high even for $F = 4$. For larger values of F, the proposed method out-performs the DCT method. As shown in Fig. 11(b), the resulting image for $F = 7$ shows visible degradation due to the resolution adjustment, but the watermark can still be detected for the proposed scheme.

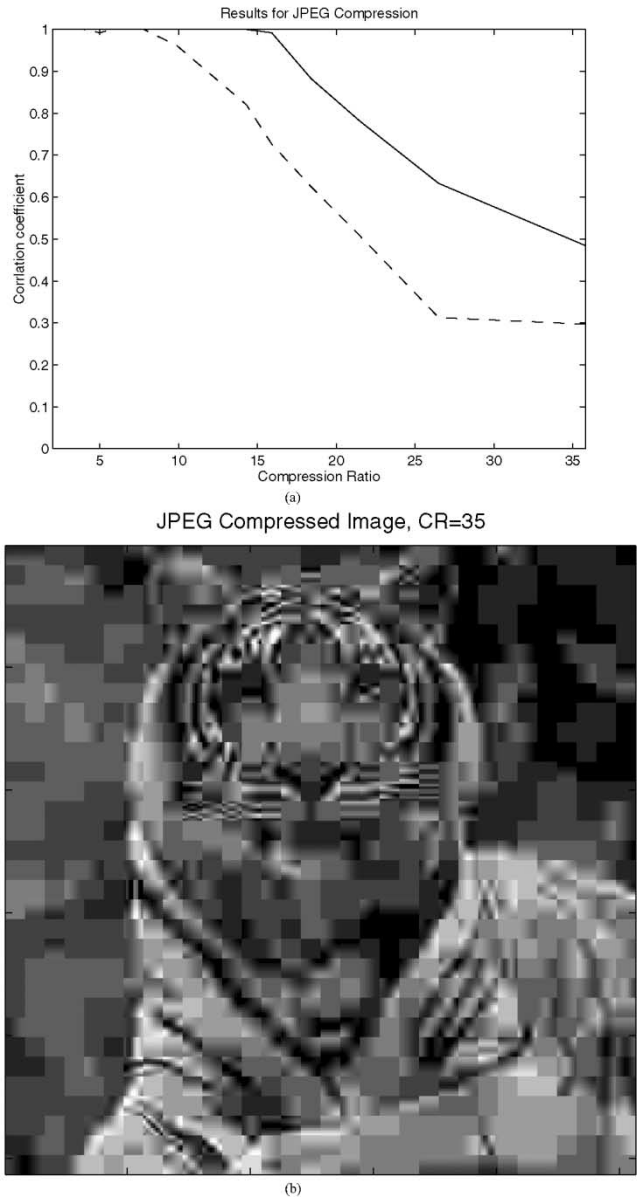

Fig. 9. Results for JPEG compression: (a) correlation coefficient versus compression ratio (CR) (the dashed and solid lines correspond to the DCT and the proposed methods, respectively) and (b) degraded watermarked image for a CR of 35.

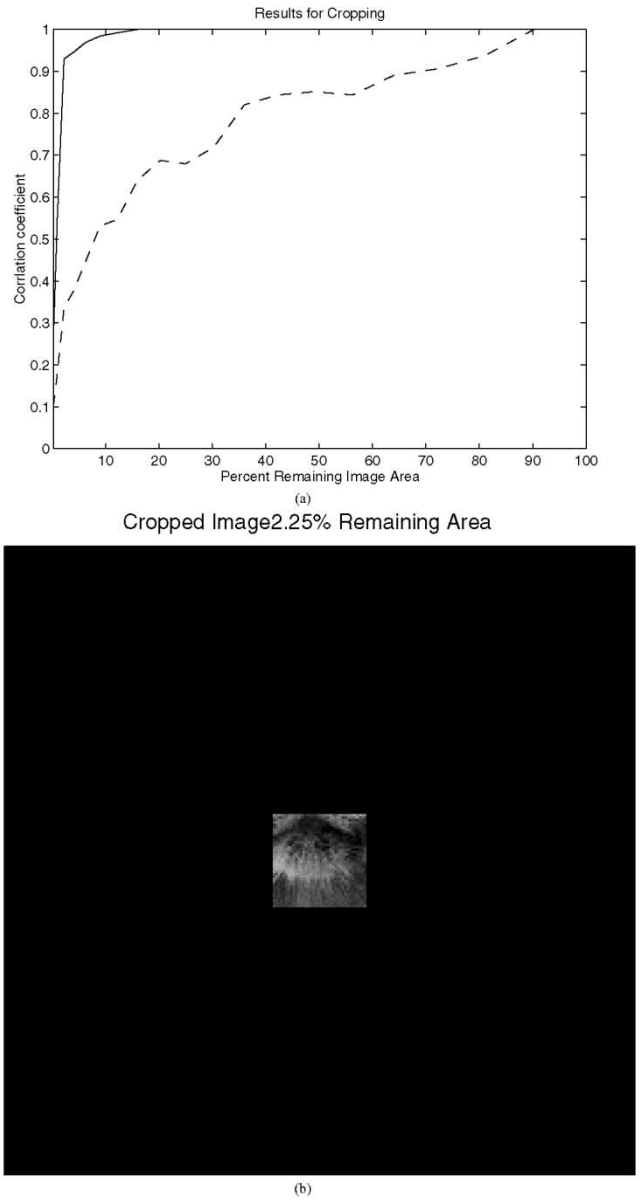

Fig. 10. Results for cropping: (a) correlation coefficient versus percent remaining image area (the dashed and solid lines correspond to the DCT and the proposed methods, respectively) and (b) cropped image consisting of 2.25 percent of the original image area.

B. Results for Logo Watermarking

We compare the results of our proposed scheme in this paper to that of [24] previously proposed by the authors. Both techniques are nonblind and have the same embedding procedure. FuseMark, however, also incorporates minimum variance fusion for watermark estimation. In [24], simple averaging of the embedded watermark repetitions is employed. FuseMark was not compared to other logo watermarking schemes due to the incompatibility in mark characteristics and host accessibility.

The host image, watermarked image and associated logo image are shown in Fig. 12(a)–(c), respectively. The logo is an 8-bit (i.e., 256 color) 32×32 grayscale image. To form the watermark, the DC value is subtracted from the logo image and the result is scaled between -1 and 1. Watermarking was performed using the following parameters: $\overline{s} = 1.0$, $\tilde{s} = 0.1$, $L = 4$, and $B = 75$.

Degradations similar to the ones discussed in Section IV-A were applied. When the watermark was extracted it was scaled, so that its minimum pixel value was set to black and its maximum pixel value to white. The watermark estimates for various distortions are displayed in Figs. 13–16. We provide representative results for situations in which the watermarked image was significantly distorted. We do not show the corresponding distorted watermarked images as they are similar to those provided in Section IV-A for watermark detection. The weighted watermark extraction produces clearer logo estimates than simple averaging in all cases. In general the watermark logos were highly resilient to additive noise and JPEG compression. The results shown in Fig. 15 for JPEG compression

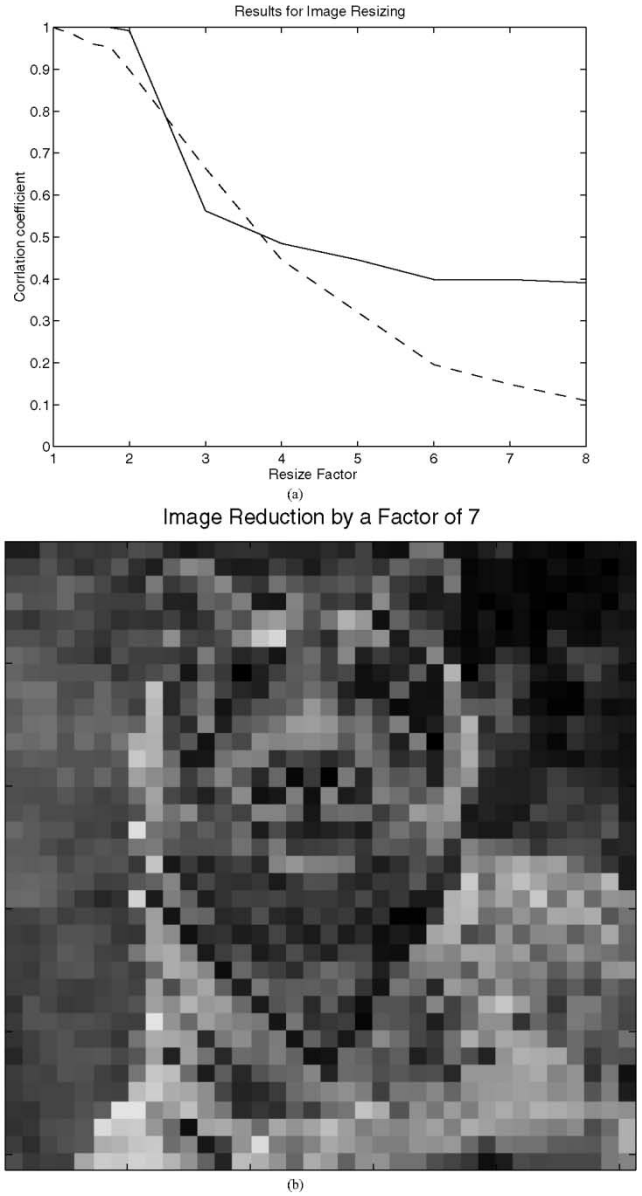

Fig. 11. Results for image resizing: (a) correlation coefficient versus image resize factor (the dashed and solid lines correspond to the DCT and the proposed methods, respectively) and (b) image size reduced by a factor of 7.

ratios of 15 and 30 correspond to compressed watermarked images which exhibit visible blocking artifacts.

V. Discussion

A. Testing and Results

In general, FuseMark performed comparably to or better than the DCT method for most images and attacks. The use of the DWT domain inherently makes our design more resilient to localized spatial and frequency domain distortions including cropping, resolution reduction and filtering. From experience with other host images, we find that the proposed method works better in general for images with highly varying localized characteristics (i.e., images with both smooth and busy areas). This is due to the fact that our HVS-based merging rule adapts the watermark signal strength to the local masking characteristics of the host image. Thus, a higher energy signal can be imperceptibly embedded within many regions of the signal. The proposed method also works better than the DCT technique for images with an overall high variance as demonstrated in Section IV-B

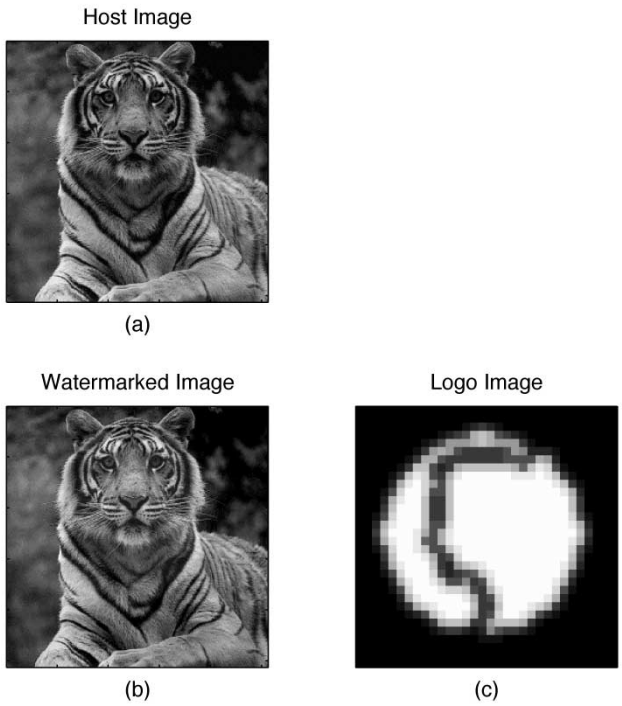

Fig. 12. Results for logo watermarking using the proposed fusion-based technique: (a) host image, (b) watermarked image, and (c) embedded 32×32 256 grayscale logo.

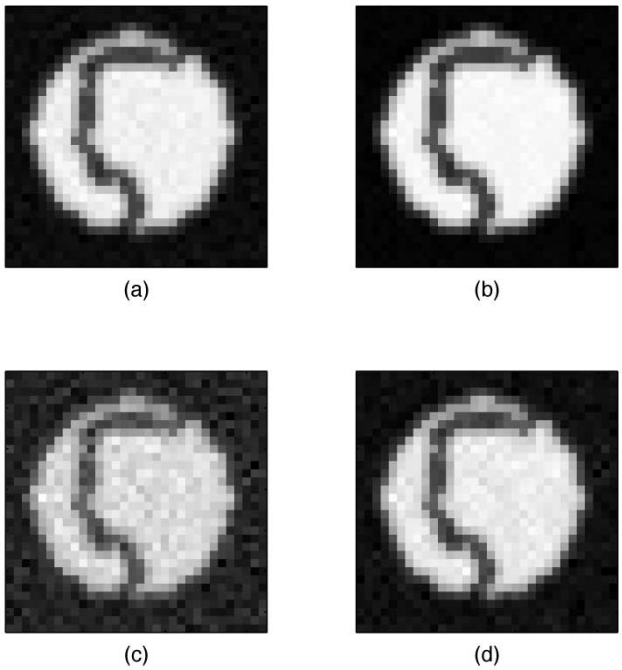

Fig. 13. Results of logo watermarking for additive noise: (a) 30 dB SNR, extraction by averaging, (b) 30 dB, SNR, extraction by optimal weighting, (c) 20 dB SNR, extraction by averaging, and (d) 20 dB, SNR, extraction by optimal weighting.

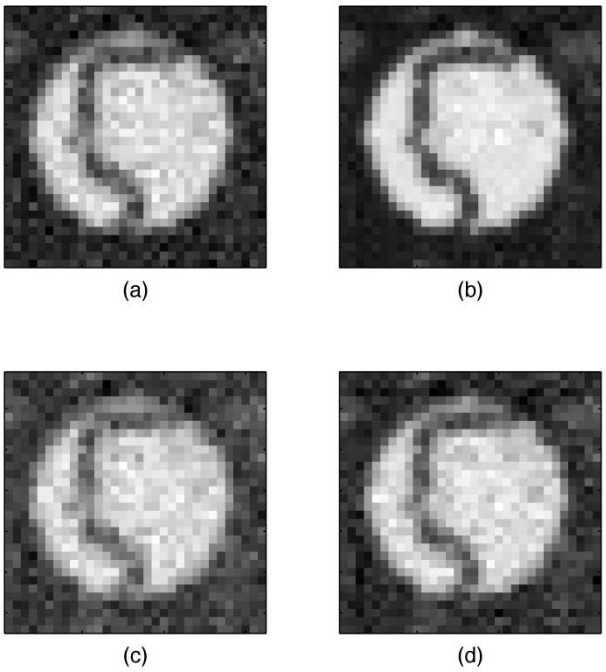

Fig. 14. Results of logo watermarking for median filtering: (a) 3×3 median filtering, extraction by averaging, (b) 3×3 median filtering, extraction by optimal weighting, (c) 5×5 median filtering, extraction by averaging, and (d) 5×5 median filtering, SNR, extraction by optimal weighting.

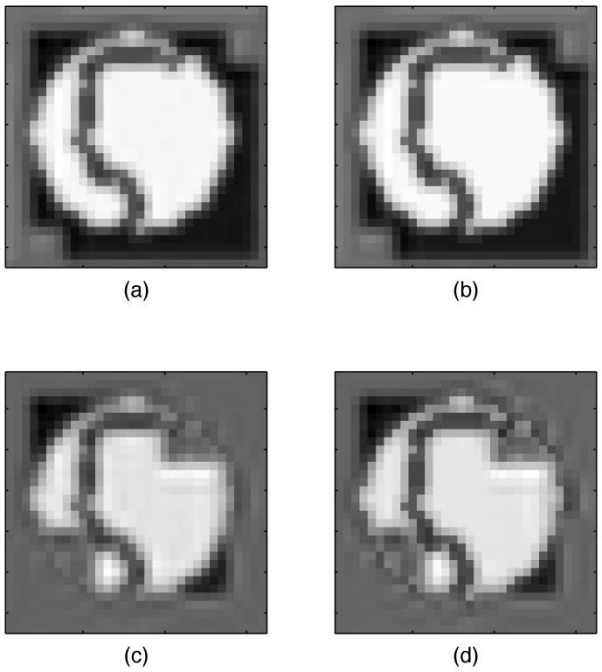

Fig. 16. Results of logo watermarking for cropping: (a) 56% of image area remaining, extraction by averaging, (b) 56% of image area remaining, extraction by optimal weighting, (c) 25% of image area remaining, extraction by averaging, and (d) 25% of image area remaining, extraction by optimal weighting.

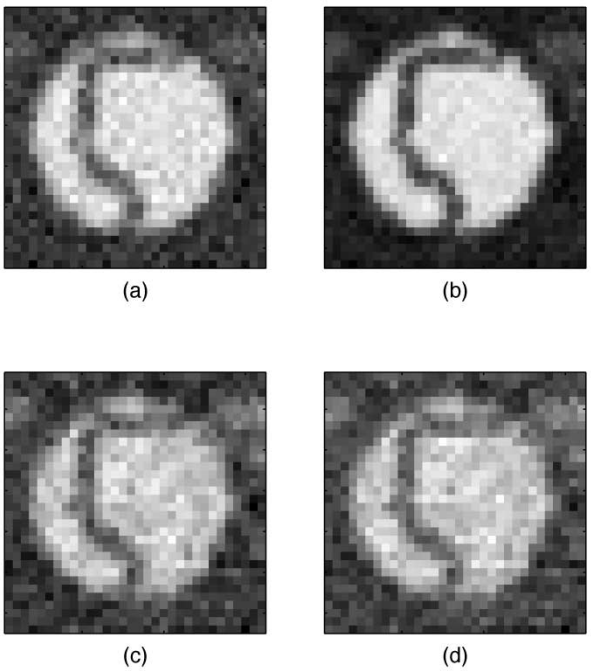

Fig. 15. Results of logo watermarking for JPEG compression: (a) CR of 15, extraction by averaging, (b) CR of 15, extraction by optimal weighting, (c) CR of 30, extraction by averaging, and (d) CR of 30, extraction by optimal weighting.

with the tiger host image. The DCT method performs better than FuseMark for mean and median filtering of smooth images. This is because although FuseMark embeds an overall higher energy signal, the DCT method can embed the watermark in lower frequency components than FuseMark while maintaining imperceptibility of the mark. One way in which to improve FuseMark would be to increase the value of the maximum decomposition level L. However, the chance of watermark visibility increases.

In addition to the results demonstrated in Section IV-B for logo marking, other types of logos such as a low resolution 32×32 image of Lena were embedded in the host image. Generally, logos in contrast with statistical watermarks can be embedded with much stronger energy while remaining imperceptible within the host signal. We found that the larger the size of the logo, the less resistance the technique had to cropping. This is due to the fact that the watermark is less spatially localized within the image due to its larger size. Experimentally, we found that the embedded logo undergoes a proportional level of *perceptible* distortion as the watermarked image. This is an inherent advantage to our fusion-based watermarking scheme since an attacker would have to essentially destroy the watermarked image to guarantee that the logo watermark was sufficiently degraded. The use of weighted watermark extraction significantly improved the quality of the extracted watermark for AWGN degradation, mild filtering and JPEG compression.

Although the distortions that we have tested are not exhaustive, they provide an indication of the potential of the proposed technique and its relative performance with the DCT method. Combinations of attacks were simulated as well. In most cases, the quality of the host signal degraded more quickly than the watermark detection capability. This can be attributed to the weighted watermark extraction technique that works well on general distortions as well as knowledge of the host signal which allows for a sophisticated embedding procedures using salience measures. The DCT method works better than the proposed technique for combined distortions which only target high and middle frequency components.

B. Geometric Distortions

Our watermark degradation model used for extraction does not account for de-synchronization attacks. Since FuseMark is not blind, one way to overcome this obstacle is to try and register the manipulated watermarked image with the host signal. For arbitrary nonlinear distortions this procedure may be tricky. However, the authors believe that for affine-like transformations, the accessibility to the host is a great advantage.

Image fusion strategies applied in this work are employed for both data embedding and extraction. Since most fusion algorithms assume that registration of images takes place prior to the merging procedure, it is not surprising that proper registration is required in our scheme for accurate watermark estimation.

C. False Detection Statistics

The analytic results presented in Section III serve to highlight a limitation of detection strategies that are directly dependent on the watermark to be detected. To decrease the probability of false negative detection, the search space of the watermark is increased and a decision is made based on the "best" correlation. Unfortunately, it is often the case that if you look too hard, you may eventually see what you are searching for (even if it is not there!).

Our proposed watermark estimation, which is fundamentally a form of minimum variance fusion, increases robustness without sacrificing false positive detection. The key idea is that detection is made more robust by employing a measure of watermark reliability that is *independent* of the watermark being tested for. In our case, the host image is employed to get a measure of the "noise." In cases, that the scheme is blind, a *reference watermark* may be employed as discussed by the authors in [31]–[33]. The fundamental insight is that in a watermark extraction scheme, the false negative detection rate must be optimized with respect to an unbiased quantity independent of the mark.

D. Security

We focus on signal processing aspects to improve robustness of logo watermarking. This technology, we hope, can be useful for an overall audit, tagging or copy protection system that accounts for protocol attacks [34], blind watermark estimation and removal, and collusion. Our feeling is that modifications of the algorithm will help solve at least some of these problems.

For instance, the embedding strategy may be made host image dependent through the use of one-way cryptographic functions as discussed in [34]. In this way, it is difficult to create a counterfeit "original" host image. One preliminary idea that will be the focus of future work is to relate the embedding direction to a hashed version of the host. That is, we take the hash of the host to produce a two-dimensional sequence $h(m,n) \in \{0,1\}$ for $m = 1, 2, \ldots, N_{wx}$, $n = 1, 2, \ldots, N_{wy}$. We can then replace (3) with

$$f_{o,l}^{w,i}(m,n) = f_{o,l}^{i}(m,n) + (2h(m,n) - 1)\, s_{o,l}^{i} w_{o,1}(m,n). \tag{24}$$

Given the resulting hash is one-way, producing a fake host such that the embedding follows (24) will be difficult.

Existing logo schemes gain security against mark access by scrambling the watermark using a secret key [15], [16]. The advantage in this paper of *not* scrambling the mark is that the perceptual characteristics of the logo may be exploited to hide it more effectively. In FuseMark, the same detail components are embedded in the corresponding details of the host, so that their characteristics match up; this helps with both imperceptibility and robustness. Attacks that attempt to maintain the perceptual quality of the watermarked image are forced in some sense to do the same for the mark. Scrambling the logo before embedding or applying a noninvertible embedding rule such as in ((24) will also scramble the perceptual characteristics of the mark. Hence, there appears to be an underlying tradeoff between security and obscurity.[4]

E. Logo Watermarks for Intellectual Property Disputes

As mentioned in the introduction, FuseMark is motivated, in part, by the assertion that a logo watermark may be more understandable to nontechnical arbitrators than statistical watermarks for supporting infringement claims. There may be an analogy between our use of statistical watermarks versus watermark logos, and fingerprints versus DNA evidence in jury-based court cases. In the latter situation, DNA evidence is adopted with much more reliability than fingerprints which brings about the natural question: are statistical watermarks, therefore, more reliable than logos?

Our practical experience with both statistical and logo watermarking leads us to believe that although there is a stronger mathematical basis for watermark detection, there is not yet a clear practical advantage of using detection strategies over logos. DNA evidence has a much stronger scientific basis than fingerprint analysis, in part, because removing or modifying DNA is not possible for nonspecialists. However, watermark detection analysis heavily lies on statistical modeling which is not yet certain to be significantly more accurate than those employed for logo watermarking. However, we would like to emphasize that, overall, the success of using logo watermarks as supporting evidence in disputes is dependent, in part, on the ability of individuals to distinguish between logos which is still an open research problem to our knowledge.

F. Grayscale Versus Binary Logos

The fundamental advantage that we exploit in this work is based on the characterisitics of grayscale logos that are more easily masked by natural images, and, hence, can be more strongly and judiciously embedded in the host for more reliable extraction. Masking theory from the area of human factors [35] demonstrates that the human perception system does not perceive as well signals that are "similar" to one another. Hence, we assert that using our basic structure for embedding, there is more success in hiding higher energy grayscale logos than binary logos in a host that is a natural image; the spectrum of the grayscale logo and host will be somewhat matched.

[4]By *obscurity* we mean data hiding capability.

VI. Conclusions

We have investigated the use of image fusion principles for the problem of robust logo watermarking; an HVS model has been incorporated to determine salient components of the image in which to embed the watermark. The watermark extraction process, that may also be viewed as a logo image refusion process, involves assessing signal attacks and appropriately weighting and recombining the embedded watermark components to maximize SNR. This strategy of weighted watermark extraction may also be applied to watermark detection algorithms in order to decrease false negative detection without increasing false positive rate.

FuseMark, in its present form, demonstrates some performance advantages of borrowing image fusion tool-sets for the problem of watermarking. Several limitations must, however, be overcome for the method to be useful for practical application. These include making the watermark recoverable for blind detection, and automating the process of determining the maximum overall watermark energy. Future work will address applying fusion principles for audio signals as well.

Acknowledgment

The authors would like to thank the associate editor, Dr. M. Yeung, and the three anonymous reviewers of this paper for their helpful comments.

References

[1] I. J. Cox, J. Killian, T. Leighton, and T. Shamoon, "Secure Spread Spectrum Watermarking for Multimedia," NEC Res. Inst., Tech. Rep. 95–10, 1995.

[2] C.-T. Hsu and J.-L. Wu, "Hidden signatures in images," in *Proc. IEEE Int. Conf. Image Processing*, vol. 3, 1996, pp. 223–226.

[3] ——, "Digital watermarking for video," in *Proc. 13th Int. Conf. Digital Signal Processing*, 1997, pp. 217–220.

[4] B. Tao and B. Dickinson, "Adaptive watermarking in the DCT domain," in *Proc. Int. Conf. Acoustics, Speech, and Signal Processing*, vol. 4, 1997, pp. 2985–2988.

[5] X.-G. Xia, C. G. Boncelet, and G. R. Arce, "A multiresolution watermark for digital images," in *Proc. IEEE Int. Conf. Image Processing*, vol. 1, 1997, pp. 548–551.

[6] ——, "Wavelet transform based watermark for digital images," *Opt. Express*, vol. 3, pp. 497–508, December 1998.

[7] C. I. Podilchuk and W. Zeng, "Image-adaptive watermarking using visual models," *IEEE Jo. Select. Areas Commun.*, vol. 16, pp. 525–539, May 1998.

[8] F. Mintzer, G. W. Braudaway, and A. E. Bell, "Opportunities for watermarking standards," *Commun. ACM*, vol. 41, pp. 57–64, July 1998.

[9] G. W. Braudaway, "Protecting publicly-available images with an invisible image watermark," in *Proc. IEEE Int. Conf. Image Processing*, vol. 1, Oct. 1997, pp. 524–527.

[10] B. L. W. Zeng and S. Lei, "Extraction of multiresolution watermark images for resolving rightful ownership," *Proc. SPIE, Security and Watermarking of Multimedia Contents*, vol. 3657, pp. 404–414, Jan. 1999.

[11] G. W. Braudaway, K. A. Magerlein, and F. Mintzer, "Protecting publically-available images with a visible image watermark," *Proc. SPIE, Optical Security and Counterfeit Deterrence Techniques*, vol. 2659, pp. 126–133, Feb. 1996.

[12] A. R. Rao, G. W. Braudaway, and F. C. Mintzer, "Automatic visible watermarking of images," *Proc. SPIE, Optical Security and Counterfeit Deterrence Techniques II*, vol. 3314, pp. 110–121, 1998.

[13] J. Meng and S.-F. Chang, "Embedding visible video watermarks in the compressed domain," in *Proc. IEEE Int. Conf. Image Processing*, vol. 1, 1998.

[14] M. S. Kankanhalli, Rajmohan, and K. R. Ramakrishnan, "Adaptive visible watermarking of images," in *Proc. IEEE Conf. Multimedia Computing and Systems*, vol. 1, 1999, pp. 568–573.

[15] G. Voyatis and I. Pitas, "Digital image watermarking using mixing systems," *Comput. Graph.*, vol. 22, no. 4, pp. 405–416, 1998.

[16] P.-L. Lin, "Robust transparent image watermarking system with spatial mechanisms," *J. Syst. Softw.*, vol. 50, pp. 107–116, 2000.

[17] P.-C. Su and C.-C. J. Kuo, "Blind digital watermarking for cartoon and map images," *Proc. SPIE, Security and Watermarking of Multimedia Contents*, vol. 3657, pp. 296–306, Jan. 1999.

[18] X. Niu, S. Sun, and W. Xiang, "Multiresolution watermarking for video based on gray-level digital watermark," *IEEE Trans. Consumer Electron.*, vol. 46, pp. 375–384, May 2000.

[19] K. T. Knox, "Reversible digital images," *Proc. SPIE, Security and Watermarking of Multimedia Contents*, vol. 3657, pp. 397–401, Jan. 1999.

[20] M. A. Abidi and R. C. Gonzalez, *Data Fusion in Robotics and Machine Intelligence*. Toronto, ON, Canada: Academic, 1992.

[21] T. A. Wilson, S. K. Rogers, and L. R. Myers, "Perceptual-based hyperspectral image fusion using multiresolution analysis," *Opt. Eng.*, vol. 34, pp. 3154–3164, Nov. 1995.

[22] L. J. Chipman, T. M. Orr, and L. N. Graham, "Wavelets and image fusion," *Proc. SPIE*, vol. 2569, pp. 208–219, 1995.

[23] D. A. Yocky, "Image merging and data fusion by means of the discrete two-dimensional wavelet transform," *J. Opt. Soc. Amer. A*, vol. 12, pp. 1834–1841, September 1995.

[24] D. Kundur and D. Hatzinakos, "A robust digital image watermarking method using wavelet-based fusion," in *Proc. IEEE Int. Conf. Image Processing*, vol. 1, 1997, pp. 544–547.

[25] M. D. Levine, *Vision in Man and Machine*. New York: McGraw-Hill, 1985.

[26] D. Kundur, "Multiresolution Digital Watermarking: Algorithms and Implications for Multimedia Signals," Ph.D. dissertation, Univ. Toronto, Toronto, ON, Canada, 1999.

[27] R. Price and P. E. Green Jr, "A communication technique for multipath channels," *Proc. IRE*, vol. 46, pp. 555–570, Mar. 1958.

[28] M. K. Simon, J. K. Omura, R. A. Scholtz, and B. K. Levitt, *Spread Spectrum Communications Handbook*. Toronto, ON, Canada: McGraw-Hill, 1994.

[29] M. L. Miller, I. J. Cox, and J. A. Bloom, "Watermarking in the real world: An application to DVD," in *Proc. Multimedia and Security Workshop at ACM Multimedia '98*, 1998, pp. 71–79.

[30] I. Daubechies, "Orthonormal bases of compactly supported wavelets," *Commun. Pure Appl. Math.*, vol. 41, pp. 909–996, Nov. 1988.

[31] D. Kundur and D. Hatzinakos, "Improved robust watermarking through attack characterization," *Opt. Express*, vol. 3, pp. 485–490, Dec. 1998.

[32] ——, "Attack characterization for effective watermarking," presented at the IEEE Int. Conf. Image Processing, Oct. 1999.

[33] ——, "Diversity and attack characterization for improved robust watermarking," *IEEE Trans. Signal Processing*, vol. 29, pp. 2383–2396, Oct. 2001.

[34] S. Craver, N. Memon, B. Yeo, and M. M. Yeung, "Resolving rightful ownerships with invisible watermarking techniques: Limitations, attacks, and implications," *IEEE J. Select. Areas Commun.*, vol. 16, pp. 573–586, May 1998.

[35] M. S. Sanders and E. J. McCormick, *Human Factors in Engineering and Design*, 7th ed. New York: McGraw-Hill, 1993.

Deepa Kundur (S'93–M'99–SM'03) was born in Toronto, ON, Canada. She received the B.A.Sc., M.A.Sc., and Ph.D. degrees, all in electrical and computer engineering, in 1993, 1995, and 1999, respectively, from the University of Toronto.

In January 2003, she joined the Electrical Engineering Department at Texas A&M University, College Station, where she is a member of the Wireless Communications Laboratory and holds the position of Assistant Professor. From September 1999 to December 2002, she was an Assistant Professor at the Edward S. Rogers Sr. Department of Electrical and Computer Engineering, University of Toronto, where she held the title of Bell Canada Junior Chair-holder in Multimedia. Her research interests include multimedia and network security, video cryptography, data hiding and steganography, covert communications, and nonlinear and adaptive information processing algorithms.

Dr. Kundur was recently the recipient of the 2002 Gordon Slemon Teaching of Design Award and the 2002 Best Electrical Engineering Professor Award (Spring) presented by the ECE Club at the University of Toronto. She has been on numerous technical program committees and has given tutorials at ICME 2003 and Globecom 2003 in the area of digital rights management. She is a Guest Editor for the PROCEEDINGS OF THE IEEE Special Issue on Enabling Technologies for Digital Rights Management.

Dimitrios Hatzinakos (M'90–SM'98) received the Diploma degree from the University of Thessaloniki, Greece, in 1983, the M.A.Sc. degree from the University of Ottawa, Ottawa, ON, Canada, in 1986, and the Ph.D. degree from Northeastern University, Boston, MA, in 1990, all in electrical engineering.

In September 1990, he joined the Department of Electrical and Computer Engineering, University of Toronto, Toronto, ON, Canada, where he is now Professor with tenure. Also, he has served Chair of the Communications Group of the Department since July 1, 1999. His research interests are in the areas of digital communications and signal processing with applications to wireless communications, image processing and multimedia. He has organized and taught many short courses on modern signal processing frameworks and applications devoted to continuing engineering education and given numerous seminars in the area of blind signal deconvolution. He is author/coauthor of more than 120 papers in technical journals and conference proceedings and he has contributed to seven books in his areas of interest. His experience includes consulting through Electrical Engineering Consociates, Ltd. and contracts with United Signals and Systems, Inc., Burns and Fry, Ltd., Pipetronix, Ltd., Defense Research Establishment Ottawa (DREO), Vaytek, Inc., Nortel Networks, and Vivosonic, Inc. He has served as Associate Editor for the IEEE TRANSACTIONS ON SIGNAL PROCESSING (1998–2000) and the Guest Editor for Elsevier's *Signal Processing* special issue on Signal Processing Technologies for Short Burst Wireless Communications (October 2000). He was a member of the IEEE Statistical Signal and Array Processing Technical Committee (SSAP) from 1992 until 1995 and Technical Program co-Chair of the 5th Workshop on Higher-Order Statistics in July 1997. He is a member of EURASIP, the Professional Engineers of Ontario (PEO), and the Technical Chamber of Greece.