Watermark Embedding: Hiding a Signal Within a Cover Image

Mauro Barni, University of Siena
Christine I. Podilchuk, Bell Laboratories
Franco Bartolini, University of Florence
Edward J. Delp, Purdue University

ABSTRACT

When looked at as a communication task, the watermarking process can be split into three main steps: watermark generation and embedding (information transmission), possible attacks (transmission through the channel), and watermark retrieval (information decoding at the receiver side). In this article we review the main issues in watermark generation and embedding. By focusing on the case of image watermarking, we first discuss the choice of the image features the watermark is superimposed to. Then we consider watermark generation and the rule used to insert the watermark within the host features. By adopting again a communication perspective, some useful hints are given on the way the watermark should be shaped and inserted within the host document for increased robustness against attacks. Given that invisibility is one of the main requirements a watermark must satisfy, the way psycho-visual notions can be used to effectively hide the watermark within an image is carefully reviewed in the second part of the article. Rather than insisting on the mathematical aspects of each of the above issues, the main rationale behind the most commonly adopted approaches is given, as well as some illustrative examples.

INTRODUCTION

Among the possible approaches to the protection of copyrighted data, digital watermarking is receiving increasing attention, as it represents a viable solution to data protection in open, highly uncontrolled, environments where cryptographic techniques cannot be applied successfully [1]. Additionally, watermarking technology is being addressed in different application scenarios, such as data authentication, database indexing, error recovery, or audio/video resynchronization

According to the watermarking approach, protection is achieved by embedding a piece of information, i.e. the watermark, within the to-be-protected data (host or cover document). Generally, the embedded watermark must be imperceptible so that the quality of the document is not affected by the presence of the watermark. At any given moment the embedded information can be extracted to prove ownership, to ensure integrity, or simply to get some copyright-related information. Depending on the application, the watermark is requested to survive all the possible manipulations the host data may undergo, with the only constraint that manipulations must not degrade *too much* the quality of the document. This is the case, for example, of copyright protection and ownership verification applications.

The first step in the design of a watermarking system is the definition of the embedding procedure. This is a crucial task, since watermark properties highly depend on the way the watermark is inserted within the data. From a very general point of view, watermark embedding is achieved by first extracting a set of features (host features) from the host data, and then by modifying them according to the watermark content. Thus, two steps are required in order to define the embedding process: choice of host features, and definition of the embedding rule. Several solutions have been proposed, leading to different classes of watermarking systems. In this article we will review the main approaches proposed so far, paying attention to discuss the advantages and the drawbacks of systems operating in different feature domains and adopting different embedding rules. We will only consider the image watermarking case, partly because most of the research developed so far focuses on the image case, and partly because many of the concepts we will discuss can be easily extended to the watermarking of different media. More specifically, we will first discuss the choice of

0163-6804/01/$10.00 © 2001 IEEE

the host features, and its implication on watermark robustness and imperceptibility, then we will address the definition of the embedding rule.

The interference between the host data and the watermark signal plays a crucial role in the design of a robust watermarking scheme, especially in blind systems, where watermark recovery is performed without any reference to the original, non-marked image. Early systems modeled the host data as a disturbing noise limiting the effectiveness of watermark communication. However, a more accurate analysis reveals that it is possible to compensate for host data interference by properly designing the embedding strategy. This is the topic of a later section, where some hints on how to embed the watermark are obtained by adopting an information theoretic point of view in the last part of this work.

The joint achievement of watermark invisibility (once again we are restricting our analysis to the image case) and robustness requires that the main properties of the Human Visual System (HVS) are exploited. Watermark robustness, in fact, calls for the watermark to be as strong as possible, a requirement that obviously conflicts with the invisibility constraint. Stated in another way, it is mandatory that the characteristics of the HVS are taken into account to embed a watermark that can hardly be perceived by the human eye, otherwise too weak a watermark would be inserted due to the invisibility requirement. The exploitation of HVS properties can be pursued either implicitly, by properly choosing the embedding domain and the embedding strategy, or explicitly, by inserting an ad hoc visual masking module that is in charge of reshaping the watermark content according to HVS considerations. The exploitation of HVS characteristics for improved watermark concealment is discussed in the last part of this work.

Choice of Host Features

In designing an effective watermarking system, it is important to determine the host feature set for embedding the watermark information. Many watermarking applications require a scheme whereby the watermark modifications do not alter the perceptual quality of the host signal. In other words, the watermarked host signal should be identical to the unwatermarked host signal in terms of visibility, audibility, intelligibility, or some other relevant perceptual criterion. Another important requirement for effective watermarking is robustness to signal processing alterations that intentionally or unintentionally attempt to remove or alter the watermark information. The feature set and embedding rules should provide a watermark that is difficult to remove or alter without severely degrading the integrity of the original host signal. For other applications, capacity rather than robustness may be a critical component. Here, capacity refers to the payload or the amount of watermark information that can be reliably hidden and recovered with low probability of error. Depending on the application, different watermarking schemes have been proposed that address to various degrees some or all of the requirements of imperceptibility, robustness, and capacity along with other issues such as cost, complexity, and whether the original host signal is available for watermark detection.

Watermark embedding can be applied directly to the original signal space of the host document or in some transform domain in order to exploit perceptual properties and/or robustness to certain signal processing transformations. For example, direct embedding of the watermark signal can be applied to the pixel values of a digital image. Typical representations include an 8-bit grayscale or 24-bit color representation for each pixel value in the image. Many times, direct embedding in the original signal space is desirable for low complexity, low cost, low delay, or some other system requirements. The embedding rule determines the pixel locations and strength of the watermark signal to be embedded and will be discussed in more detail later. The location for watermark embedding can be determined by low-level waveform processing or some higher-level processing such as edge detection or feature extraction.

Transform domain embedding for digital image watermarking includes transforms such as the block-based discrete cosine transform (DCT), the discrete wavelet transform (DWT), as well as other frequency domain representations. The block-based discrete cosine transform is a popular choice for watermarking image data because it is a basic component of image and video compression standards such as JPEG and the MPEG and ITU H.26x family of coders. By choosing a framework that matches current compression standards, watermark embedding schemes can be designed to avoid adding watermark information to the coefficients that are typically discarded or coarsely quantized, resulting in a scheme that is robust to compression. Another important reason to choose transform domain watermarking that matches the framework of current compression standards is that for many applications, direct embedding in a compressed bitstream is a desirable or necessary feature. This is especially true for some video watermarking applications where the video will most likely be in some compressed form such as an MPEG2 bitstream, and it is desirable to add the watermark information directly to the MPEG2 bitstream with only a partial decode.

Perceptual design constraints that guarantee invisibility can also be readily incorporated into frequency domain representations, by avoiding watermarking low frequency components where alterations may produce very visible distortions. Robustness issues can also be addressed by choosing transform domains for watermark embedding that are invariant to certain types of transformations. For instance, applying a watermark in the Fourier-Mellin domain results in a watermark that is invariant to image translation, scale, and rotation. Applying a watermark to the DWT coefficients of an image results in a multiresolution watermark signal where the scale of the watermark makes it robust to different types of alterations. Perceptual factors are also easily incorporated into a wavelet-based watermarking algorithm.

The watermark information can be embedded directly in an image, video, or audio signal by altering the signal values either in the original spatio-temporal domains or in some transform

Transform domain embedding for digital image watermarking includes transforms such as the block-based discrete cosine transform (DCT), discrete wavelet transform (DWT), as well as other frequency domain representations.

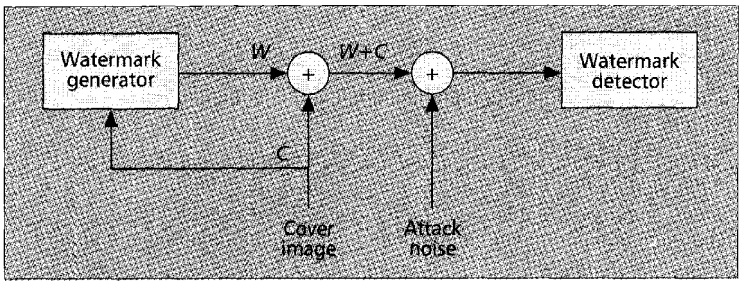

■ **Figure 1.** *Watermark embedding and recovering according to the informed-embedding, blind-decoding paradigm.*

domain. For other types of content, such as a postscript file of an electronic version of a text document, the watermark information can be cleverly embedded into the actual format of the document, that is, by altering the spacing between lines, between characters and words or by minor alterations to the characters themselves. Similarly, graphical representations such as the parameters used for facial animation as defined by MPEG-4 can be altered slightly to embed additional information without noticeably distorting the facial features associated with these parameters.

The next section addresses the general embedding rules that have been proposed for embedding the watermark information into the host features described here.

Embedding Rule

Once the embedding domain has been chosen, the rule used to blend the watermark and the host features together must be defined. Prior to defining the embedding rule, however, it is worthwhile to examine the watermark shape, since the ultimate performance of the watermarking system also depends on the form of the watermark signal.

Watermark Shaping

In most of the systems proposed so far, the watermark consists of a pseudo-random sequence of independent and identically distributed samples. Such a form derives from a communications approach to the watermarking problem, where watermarking is looked at as the transmission of a weak signal over a very noisy channel, a problem that is commonly tackled by means of spread spectrum techniques. The pseudo-random sequence is usually generated by starting from a secret key to achieve system security. Then, the sequence is treated itself as the watermark or it is modulated by an antipodal bit sequence. In the former case the decoder is only asked to decide upon the watermark presence (1-bit watermark), whereas in the latter case the modulating bits are recovered through conventional spread spectrum decoding.

In some applications it is convenient that the watermark corresponds to an image containing copyright information, e.g. a visual logo or a serial number. In such a case, a key-dependent scrambling function is usually applied to the watermark before embedding, so that the signal actually embedded within the cover image resembles a pseudo-noise sequence.

A somewhat different approach results when the peculiarities of watermarking applications are taken into account in the modeling of the transmission channel. The attacker, in fact, has to satisfy some general requirements on the distortion he/she can introduce, since the attack strength is limited by the need not to degrade too much the image quality. When such a constraint is taken into account, it turns out that the watermark signal must be as similar as possible to the cover image, since in this case, it is more difficult for the attacker to distinguish between the watermark he/she wants to destroy and the image he/she wants to preserve. The exact formulation of the above principle depends on the distortion metric used to judge image quality. For instance, if the MSE criterion is used, some conditions on the power spectrum of the watermark can be obtained. Shaping of the watermark power spectrum can be performed either by filtering an intermediate pseudo-random, white watermark or by using chaotic sequences. A more interesting approach should consider the perceptibility of the degradation to a human observer instead of the MSE criterion.

Additive, Non-Additive, and Substitution Watermarks

The two most common approaches to watermark embedding are the *additive* one, for which

$$y_i = x_i + \gamma m_i \,, \tag{1}$$

where x_i is the i-th component of the original feature vector, m_i the i-th sample of the watermark, γ is a parameter controlling the watermark strength, and y_i is the i-th component of the watermarked feature vector; and the *multiplicative* one, for which

$$y_i = x_i + \gamma m_i x_i \,, \tag{2}$$

where the symbols have the same meaning as in equation (1).

The main reason for the popularity of additive watermarking is its simplicity. Additive watermarks are mainly used in the spatial domain, since in this case watermark concealment is achieved very simply by adapting the watermark strength γ to the local characteristics of the cover image. Another advantage of additive watermarking is that under the assumption that the host features follow a Gaussian distribution and that attacks are limited to the addition of white Gaussian noise (AWGN model), correlation-based decoding is optimum, in that either the overall error probability, or the probability of missing the watermark given a false detection rate, can be minimized. The adoption of correlation decoding, in turn, makes it possible to cope with spatial shifts due, for example, to image cropping. The exhaustive search of the watermark by looking at all possible spatial locations, in fact, can be accomplished efficiently in the transformed domain, since signal correlation in the spatial domain corresponds to a multiplication in the Fourier domain.

Techniques operating in the full-frame frequency domain, be it the DCT or the DFT domain, tend to adopt a multiplicative embedding rule. The main reason for such a choice lies in the masking properties of the Human Visual System. It is known, in fact, that it is more diffi-

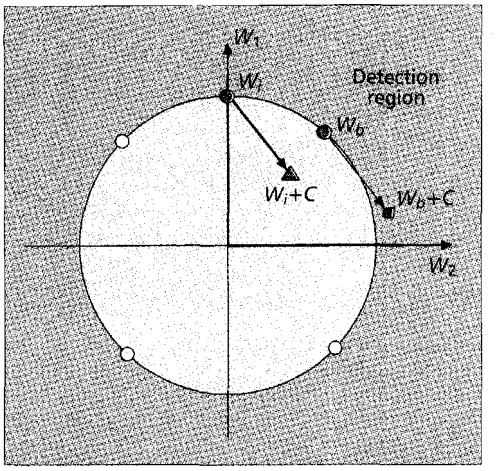

Figure 2. *Example of informed embedding. Since the encoder knows in advance the noise component due to the cover image (**C**), it embeds the watermark $\mathbf{w}_i$ instead of $\mathbf{w}_b$, thus resulting in higher robustness against noise due to attacks.*

cult to perceive a disturbance at a given frequency if the image already contains such a frequency component. In other words, for a better match of the invisibility constraint, it is preferable to embed a watermark whose energy at a given frequency is proportional to the energy of the image at that frequency. Another advantage of multiplicative watermarking is that an image-dependent watermark is obtained, thus increasing system security, since in this case it is more difficult to estimate the watermark by averaging a set of watermarked images.

Sometimes a third category of watermarking algorithms is introduced to refer to systems that operate by substituting a subset of the host features with new values, even if, truly speaking, the distinction between additive and substitutive algorithms tends to be rather vague, and some methods could be classified either as additive or substitutive. This is the case with many of the algorithms operating in the block-DCT domain, where some of the DCT coefficients are modified according to the watermark content. Possible solutions include re-quantization of DCT coefficients, substitution of coefficients, and modification of coefficients so that a given relationship is imposed on their order.

Of course, the three categories described above cannot accommodate the huge variety of watermarking algorithms proposed so far. For example, systems designed to operate on particular kinds of data contents, such as postscript text files, would deserve a separate treatment, since the embedding strategies used in that case are completely different from those usually adopted for the watermarking of natural images.

Informed Embedding

A problem with blind watermarking is that the decoder does not know the original cover image. In other words, for the decoder the original image is nothing else than noise added to the true signal, i.e. the watermark. Actually, recent research [2] demonstrated that the host image should not be treated as conventional noise, since the encoder knows it in advance (Fig. 1), and hence it can take some proper countermeasures to reduce the impact of decoder blindness on watermarking reliability. Such a watermarking strategy is usually referred to as informed watermark embedding.

By grounding on a solid information theoretic framework developed about 20 years ago [3], some very useful hints can be obtained on how to embed the watermark so that its robustness is augmented. By considering again the scheme reported in Fig. 1, we can express the general rule of informed embedding as follows: *the encoder looks at the surroundings of vector* **c** *representing the cover image, then it chooses a watermark* **w** *which is compatible with the visibility constraint and far enough from the non-detection region to be distinguishable when viewed at the decoder side.* Note that in so doing, the encoder adapts watermark generation/embedding to the state **c** (the cover image in our case) of the channel.

In an attempt to clarify, and maybe oversimplify, the informed embedding concept, let us consider the example reported in Fig. 2. The inside of the circle represents all the possible watermarks satisfying the visibility constraint (here reduced to a simpler constraint on watermark power), and point $\mathbf{w}_b$ a blind-embedding watermark, i.e., a watermark that is chosen regardless of the noise introduced by the channel. Also assume that watermark embedding follows an additive rule. When the watermark is added to the image **c**, or, to better follow the communication paradigm, when the image **c** is added to the watermark, the watermark moves to $\mathbf{w}_b + \mathbf{c}$. It is on such a new signal that the noise due to attacks operates. As it can be readily seen, the presence of the cover image contributes to the weakening of watermark robustness, since $\mathbf{w}_b + \mathbf{c}$ is closer to the non-detection region than $\mathbf{w}_b$.

The behavior of an informed-embedder would be drastically different. Such an embedder, in fact, exploits the knowledge about the noise vector **c** (the cover image), and decides to transmit a watermark that is in position $\mathbf{w}_i$. After addition of the cover image, we obtain a signal that is well within the detection region, thus resulting in a watermark that is far more robust than $\mathbf{w}_b$. Stated in another way, since the embedder knows that noise will enforce the first component of the watermark, it decides to decrease such a component to re-enforce the second one, which he knows will be severely affected by the channel.

Interestingly, and somewhat surprisingly, it can be demonstrated [3] that decoder blindness does not affect channel capacity at all, thus supporting the idea that, at least asymptotically, no loss of robustness should be expected by denying the decoder the access to the original, non-marked image (channel status).

Psychovisual Factors

The previous section reviewed the general principles behind the watermark embedding process. As mentioned, there are three basic approaches: additive embedding, multiplicative embedding,

Techniques operating in the full-frame frequency domain, be it the DCT or the DFT domain, tend to adopt a multiplicative embedding rule. The main reason for such a choice lies in the masking properties of the Human Visual System.

Additional compression gains have been realized by using more sophisticated techniques for measuring perceptual quality and applying this to the design of the compression algorithm.

and substitution. For most watermarking applications, a critical feature of the watermark embedding algorithm is the ability to provide a transparent watermark that does not noticeably alter the perceived quality of the content and is maximally robust to attack. With this in mind, an effective watermarking scheme either explicitly or implicitly applies perceptual knowledge in the embedding process.

Selecting features or the transform space for watermark embedding is often based on perceptual knowledge and choosing a space where we can decouple perceptually significant and insignificant components of the original host signal. In order for the watermark to be transparent, we wish to mark the perceptually insignificant portion of the signal. In order for the watermark to be robust to intentional and unintentional attack, we wish to mark the perceptually significant portion of the signal. Perceptual models help us design watermark embedding schemes that allow us to balance these two opposing requirements. For digital image watermarking, feature selection can occur in the spatial domain, where pixel locations or local spatial areas to be marked are chosen based on some perceptual criterion. An example of spatial domain perceptually based watermarking includes a measure of local activity such as calculating the variance over local blocks and choosing blocks of data to be marked whose variance exceeds an empirically determined threshold level. This ensures that relatively smooth blocks, where alterations may be visible, are not altered. Other approaches include embedding a textured pattern as the watermark information into the image at a location with similar texture.

A transform domain framework is ideal for applying certain properties of perception into the watermark embedding process. For instance, a common framework for watermarking digital images is the block-based discrete cosine transform (DCT) with block size 8x8 or a discrete cosine or Fourier transform over the entire image. Psychovisual properties can be used to select the transform coefficients to be marked. For instance, many watermarking schemes do not mark the low frequency components because alterations may produce a noticeable visual distortion. Typically, high frequency components are not marked as well, since removing these components will wipe out the watermark signal without introducing noticeable distortions to the original image. Also, most compression algorithms will typically discard or coarsely quantize high frequency components so that watermarking schemes that are robust to compression may choose to avoid marking these coefficients. In many cases, a fixed set or random subset of midrange frequencies are chosen for watermarking.

More sophisticated perceptual models can also be used to determine the maximum strength of the watermark signal that can be tolerated at every pixel or transform coefficient without producing visible distortions. The watermark signal strength can be determined for each feature location, as an average value for the entire image or some location in space, time, or frequency. For instance, it is common to use the magnitude of the transform domain coefficients to determine the strength of the watermark signal for that coefficient. Many techniques propose to adapt the watermark signal strength as a percentage of the host signal strength at the embedding location so that a stronger host signal corresponds to a stronger watermark signal. Using local or global image characteristics to determine the strength of the watermark signal on a fine or coarse level results in an image-adaptive scheme. Therefore, for images that are very smooth, with few details and texture, the watermark signal strength will be weaker than for highly detailed images with complex textures where a stronger watermark signal can be more effectively hidden.

The fundamental work on understanding human vision has been successfully applied to practical problems such as data compression. The goal of data compression is to represent the original digital content in a compact form for storage or transmission purposes. The goal is to minimize distortion to the original content for a target bitrate or to minimize the bitrate given a target acceptable distortion level. A meaningful distortion metric should ideally be highly correlated with the perceived quality of the content as viewed by a human observer. For this reason, it is very useful to apply knowledge about psychovisual phenomena in designing effective data compression algorithms. The traditional approaches that have been very successful for data compression are focused on removing signal redundancies and using mean square error as a way to measure perceptual quality and overall compression performance. However, additional compression gains have been realized by using more sophisticated techniques for measuring perceptual quality and applying this to the design of the compression algorithm. Many compression schemes either explicitly or implicitly through empirical design have incorporated some notion of frequency sensitivity into the compression algorithm. Frequency sensitivity refers to the visual system's sensitivity to sine wave gratings at various frequencies and is sometimes referred to as the *modulation transfer function* (MTF) of the human visual system. Frequency sensitivity is independent of image characteristics and is only a function of viewing conditions. Other properties that have been observed in describing psychovisual phenomena include image-dependent characteristics such as luminance sensitivity and contrast masking. Luminance sensitivity is the ability to detect noise against different average luminance levels; contrast masking is the ability to detect one signal in the presence of another signal and takes into account characteristics of texture and high frequency details.

Commonly, vision scientists use the terminology of *just noticeable difference* (JND) as a way of mapping visual models into a quantity that can be readily used by engineers designing signal processing algorithms for data compression. The JND thresholds are usually determined for a particular viewing condition and these values can be used to determine how much distortion can be tolerated at every location within the image subject to the imperceptibility constraint. JND values can be calculated using the previously

mentioned properties of frequency sensitivity, luminance sensitivity, and contrast masking. The JNDs provide a way of determining the maximum amount of quantization noise that can be tolerated without affecting image quality under some predefined viewing conditions. Viewing conditions could include viewing distance, image size, monitor type, and lighting conditions. This provides a direct way of mapping JNDs into quantization step sizes so that the minimum bitrate for zero distortion can be achieved. The JNDs are also ideally suited for the watermark embedding problem, where the thresholds provide the maximum alteration levels possible for the marking algorithm in order to guarantee imperceptibility. As long as the watermark signal does not alter any portion of the original data beyond the JND threshold, imperceptibility is guaranteed. Actually, the perceptual models may be more effective for watermarking than for compression where algorithmic constraints and overhead costs may prohibit the use of a fine-scale image-adaptive perceptual model. For instance, the international still image compression standard, JPEG, only allows for one quantization matrix for the entire image.

Image-adaptive watermarking algorithms that use JNDs to modulate the watermark signal strength in a DCT and wavelet-framework have been proposed [4]. Figure 3 illustrates watermarked image examples using the DCT-based image-adaptive watermarking scheme where the watermark has been embedded using different viewing conditions. The images on the left show the watermarked images while the images on the right show the actual watermark signal. The top image illustrates the watermark strength at the minimum viewing distance, and for each consecutive image the viewing distance is increased by one image height. The viewing distance in the perceptual model can be adjusted to trade off imperceptibility with watermark strength (and to some degree, robustness). Note that the structure of the watermark signal is highly correlated to the original picture, with a stronger signal in areas of high texture and details, where it is less visible. The perceptual model used here was originally developed to improve JPEG coding performance and applied to the watermarking of still images [4].

Note that for most perceptually-based watermarking schemes that adapt the strength of the watermark using perceptual information, the scaling factor is a function of the image data, so that the embedding scheme is not additive. When comparing the watermark embedding process directly to the notion of data compression using visual models, it is easy to see that in the ideal case, the optimum data compression algorithm should remove all the perceptually irrelevant data so that any watermarking scheme that provides an "invisible" mark should also be wiped out by such a scheme. In other words, an optimal compression scheme should be able to detect that the changes made by adding the watermark signal fall below the *just noticeable difference* level, and these modifications could be removed for data compression purposes. However, current compression schemes are limited in how much local adaptation is possible without making the side information needed to decode such a scheme prohibitively large. This allows for watermarking algorithms to survive state-of-the-art compression schemes. Of course, besides effective compression schemes, there are other ways to effectively remove the watermark or cause detection failure. These alternative methods are discussed in other articles of this special issue.

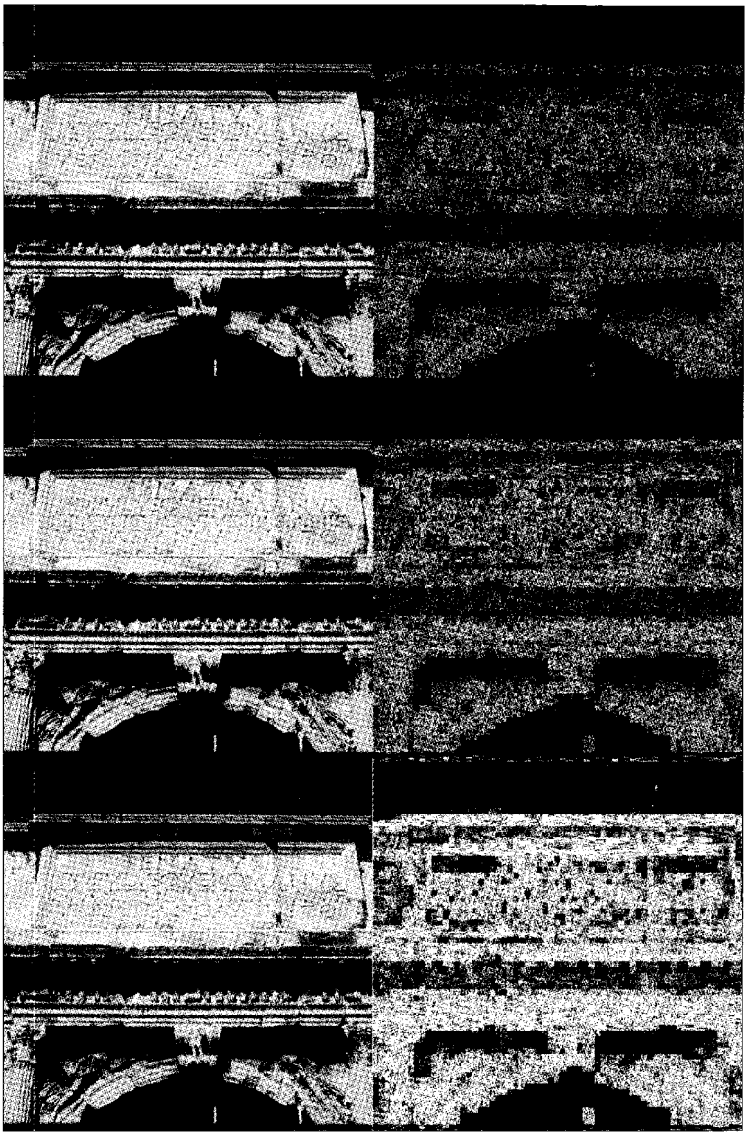

Figure 3. *Watermark example for different viewing distances.*

CONCLUSIONS

In this article the first phase of any watermarking system, i.e. watermark embedding, has been overviewed. Three issues related to watermark embedding have been identified: choice of host features, choice of the embedding rule, and exploitation of psychovisual factors. With regard to the choice of the host features the watermarking signal has to be embedded in, this strongly depends on the type of application the watermarking system is devised for. For example, techniques working directly in the original signal

The perceptual models may be more effective for watermarking than for compression where algorithmic constraints and overhead costs may prohibit the use of a fine-scale image-adaptive perceptual model.

space (spatial domain) are less cumbersome than transform domain techniques and thus are preferable if the constraint of real-time watermarking is of primary importance. On the other side, transform domain techniques can offer a superior resistance to a wider class of data manipulations, and some particular transforms (e.g., DWT) are more suitable for adapting to psychovisual constraints.

With regard to the embedding rule, the two most common approaches are the additive approach (the watermark is directly added to the host features) and the multiplicative approach (the watermark is added to the host features by weighting it proportionally to the host features values). While the former approach is used in conjunction with spatial domain techniques, the latter is preferred for transform domain techniques because it allows for better exploitation of psychovisual phenomena. From the point of view of robustness, the embedding rule should be designed in such a way that it adapts the spectral shape of the watermarking signal to that of the host data, given that it is assumed that an attacker will take care not to deteriorate too much the watermarked data. Important recent results from the analysis of the embedding rule have shown that all the information available about the host data and about the watermark detection function should be exploited for the embedding phase. In practice, the fact that the cover image, although unknown to the detector, is perfectly known to the embedder should be exploited, and thus should not be treated as noise by it (informed embedding).

With regard to the exploitation of psychovisual factors, this can partially be achieved by a careful selection of the host features, by considering those features exhibiting less sensitivity to the human eye. In addition, an explicit masking step is usually adopted to better adapt the watermark to the local image characteristics and to the properties of the HVS. A common effect of the above strategies is to highly correlate the watermarking signal to the structure of the watermarked visual data. Most of the results related to psychovisual phenomena obtained in the past in the field of visual data compression can be effectively transferred to watermarking applications. Actually, perceptual models can be even more effective for watermarking than for compression where algorithmic and overhead constraints limit their usability.

REFERENCES

[1] "Special Issue on Identification and Protection of Multimedia Information," *Proc. IEEE*, vol. 87, no. 7, July 1999.

[2] I. J. Cox, M. L. Miller, and A. L. McKellips, "Watermarking as Communications with Side Information," *Proc. IEEE*, vol. 87, no. 7, July 1999, pp. 1127–41.

[3] M. Costa, "Writing on Dirty Paper," *IEEE Trans. Information Theory*, vol. 29, no. 3, May 1983, pp. 439–41.

[4] R. B. Wolfgang, C. I. Podilchuk, and E. J. Delp, "Perceptual Watermarks for Digital Images and Video," *Proc. IEEE*, vol. 87, no. 7, July 1999, pp. 1108–26.

BIOGRAPHIES

Mauro Barni (barni@dii.unisi.it) was born in Prato, in 1965. He graduated in electronic engineering at the University of Florence in 1991. He received the Ph.D. in informatics and telecommunications in October 1995. From 1991 through 1998 he was with the Department of Electronic Engineering, University of Florence, Italy, where he worked as a postdoc researcher. Since September 1998, he has been with the Department of Information Engineering, of the University of Siena, Italy, where he works as assistant professor. His main interests are in the field of digital image processing and computer vision. His research activity is focused on the application of image processing techniques to copyright protection and authentication of multimedia data (digital watermarking), and to the transmission of image and video signals in error-prone, wireless, environments. He is author/co-author of more than 100 papers published in international Journals and Conference Proceedings. Mauro Barni is member of the IEEE, where he serves as member of the Multimedia Signal Processing Technical Committee (MMSP-TC).

Christine Podilchuk (chrisp@bell-labs.com) received the B.S., M.S. and Ph.D. degrees in electrical engineering from Rutgers University, New Brunswick NJ in 1984, 1986 and 1988 respectively. Since then, she has been with Bell Laboratories, Murray Hill, NJ. Her research interests are in signal processing, particularly for image and video applications. Her most recent work includes digital watermarking for *multimedia, video compression and video transmission* over wireless networks.

FRANCO BARTOLINI [M'96] graduated (cum laude) in electronic engineering from the University of Florence, Florence, Italy in 1991. In 1996 he received the Ph.D degree in informatics and telecommunications from the University of Florence. He is now a postdoctoral researcher with the University of Florence. His research interests include digital image sequence processing, still and moving image compression, non-linear filtering techniques, image protection and authentication (watermarking), image processing applications for the cultural heritage field; signal compression by neural networks, and secure communication protocols. He has published more than 90 papers on these topics in international journals and conferences. He holds two Italian patents in the field of digital watermarking. He is a member of IEEE and IAPR.

Edward J. Delp was born in Cincinnati, Ohio. He received the B.S.E.E. (cum laude) and M.S. degrees from the University of Cincinnati, and the Ph.D. degree from Purdue University. From 1980-1984, He was with the Department of Electrical and Computer Engineering at The University of Michigan, Ann Arbor, Michigan. Since August 1984, he has been with the School of Electrical and Computer Engineering and the Department of Biomedical Engineering at Purdue University, West Lafayette, Indiana, where he is a Professor of Electrical and Computer Engineering and Professor of Biomedical Engineering. His research interests include image and video compression, multimedia security, *medical imaging, communication and information theory*. He is a Fellow of the IEEE, a Fellow of the SPIE, and a Fellow of the Society for Imaging Science and Technology (IS&T). From 1998-2000 he was Chair of the Image and Multidimensional Signal Processing (IMDSP) Technical Committee of the IEEE Signal Processing Society. From 1994-1998 he was Vice-President for Publications of IS&T. In 2000 he was a Distinguished Lecturer of the IEEE Signal Processing Society. He was Co-Chair of the SPIE/IS&T Conference on Security and Watermarking of Multimedia Contents held in San Jose in 1999, 2000, and 2001. He was the General Co-Chair of the 1997 Visual Communications and Image Processing Conference (VCIP) held in San Jose. He was Program Chair of the IEEE Signal Processing Society's Ninth IMDSP Workshop held in Belize in March 1996. He was General Co-Chairman of the 1993 SPIE/IS&T Symposium on Electronic Imaging. He is the Co-Program Chair of the IEEE International Conference on Image Processing in 2003. From 1984-1991 he was a member of the editorial board of the International Journal of Cardiac Imaging. From 1991-1993, he was an Associate Editor of the IEEE *Transactions on Pattern Analysis and Machine Intelligence*. From 1992-1999 he was a member of the editorial board of the journal Pattern Recognition. From 1994-1999, he was an Associate Editor of the Journal of Electronic Imaging. From 1996-1998, he was an Associate Editor of the IEEE Transactions on Image Processing. In 1990 he received the Honeywell Award and in 1992 the D. D. Ewing Award for excellence in teaching. During the summers of 1998, 1999, and 2001 he was a Visiting Professor at the Tampere International Center for Signal Processing at the Tampere University of Technology in Finland.