PROCEEDINGS OF SPIE

SPIEDigitalLibrary.org/conference-proceedings-of-spie

Review of watermarking and the importance of perceptual modeling

Cox, Ingemar, Miller, Matt

Ingemar J. Cox, Matt L. Miller, "Review of watermarking and the importance of perceptual modeling," Proc. SPIE 3016, Human Vision and Electronic Imaging II, (3 June 1997); doi: 10.1117/12.274502

Find authenticated court documents without watermarks at docketalarm.com.

A review of watermarking and the importance of perceptual modeling

Ingemar J. Cox and Matt L. Miller NEC Research Institute 4 Independence Way Princeton, NJ 08540

ABSTRACT

A watermark embeds an imperceptible signal into data such as audio, video and images, for a variety of purposes, including captioning and copyright control. In this paper, we first outline the desirable characteristics of digital watermarks. Previous work in digital watermarking is then reviewed. Early work identified redundant properties of an image (or its encoding) that can be modified to encode watermarking information. The early emphasis was on hiding data, since the envisioned applications were not concerned with signal distortions or intentional tampering that might remove a watermark. However, as watermarks are increasingly used for purposes of copyright control, robustness to common signal transformations and resistance to tampering have become important considerations. Researchers have recently recognized the importance of perceptual modeling and the need to embed a signal in perceptually significant regions of an image, especially if the watermark is to survive lossy compression. However, this requirement conflicts with the need for the watermark to be imperceptible. Several recent approaches that address these issues are discussed.

1. INTRODUCTION

There has been significant recent interest in watermarking. This is primarily motivated by a need to provide copyright protection to digital content, such as audio, images and video. Digital representations of copyrighted material such as movies offer many advantages. However, the fact that an unlimited number of perfect copies can be illegally produced is a serious threat to the rights of content owners. Watermarking can be used for owner identification, to identify the content owner, fingerprinting, to identify the buyer of the content, for broadcast monitoring to determine royalty payments, and authentication, to determine whether the data has been altered in any manner from its original form. The latter purpose is somewhat different from those of copyright control and the characteristics thereof may be different and are therefore not discussed further here.

A number of technologies are being developed to provide protection from illegal copying. Two complimentary techniques are encryption and watermarking. Encryption protects content during the transmission of the data from the sender to receiver. However, after receipt and subsequent decryption, the data is no longer protected and is in the clear. Watermarking compliments encryption by embedding a signal directly into the data. Thus, the goal of a watermark is to always remain present in the data. It should be noted that embedded signaling or watermarking can be used for a variety of other purposes other than copyright control, but we restrict our discussion here to issues related to copyright control.

In the next section, we outline desirable properties of a watermark for copyright control, which can be quite different from watermarks for authentication purposes, for example, and explain why perceptually modeling is important to watermarking. Section 3introduces a framework in which to discuss the many different proposed watermarks that are described in Section 4.

2. PROPERTIES OF WATERMARKS

There are a number of desirable characteristics that a watermark should exhibit. These include that it be difficult to notice, robust to common distortions of the signal, resistant to malicious attempts to remove the watermark, support a sufficient data rate commensurate with the application, allow multiple watermarks to be added and that the decoder be scalable. These characteristics are discussed in more detail next.

DOCKE.

Difficult to notice The watermark should not be noticeable to the viewer nor should the watermark degrade the quality of the content. In earlier work,^{1,2} we had used the term "imperceptible", and this is certainly the ideal. However, if a signal is truely imperceptible, then perceptually-based lossy compression algorithms should, in principle, remove such a signal. Current state-of-the-art compression algorithms probably still leave room for an imperceptible signal to be inserted. This may not be true of next generation compression algorithms. Thus, to survive the next generation of lossy compression algorithms, it will probably be necessary for a watermark to be noticeable to a trained observer.

Of course, a just noticeable difference is usually observed by comparing two signals, e.g. compressed and uncompressed or watermarked and original. However, a typical observer will not be comparing two signals, so while a song may sound different from the original, the observer may have no way of knowing this and will probably be satisfied provided the difference is not displeasing.

Early work on watermarking focused almost exclusively on designing watermarks that were imperceptible and therefore often placed watermark signals in perceptually insignificant regions on the content. However, other properties of a watermark conflict with this choice.

Robustness Music, images and video signals may undergo many types of distortions. Lossy compression has already been mentioned, but many other signal transformations are also common. For example, an image might be contrast enhanced and colors might be altered somewhat, or an audio signal might have its bass frequencies amplified. In general, a watermark must be robust to transformations that include common signal distortions as well as digital-to-analog and analog-to-digital conversion and lossy compression. Moreover, for images and video, it is important that the watermark survive geometric distortions such as translation, scaling and cropping.

It has been argued^{1,2} that robustness can only be attained if the watermark is placed in perceptually significant regions of an image. This is because the image fidelity is only preserved if the perceptually significant regions of the image remain intact. Conversely, perceptually insignificant regions can be removed without affecting the image quality. Consequently, watermarks that are placed in perceptually insignificant regions will not be robust and can be easily removed. Note that robustness actually comprises two seperate issues: (1) whether or not the watermark is still present in the data after distortion and (2) whether the watermark detector can detect it. For example, watermarks inserted by many algorithms remain in the data after geometric distortions such as scaling, but the corresponding detection algorithms can only detect the watermark if the distortion is first removed. In this case, if the distortion cannot be determined and/or inverted, the detector cannot detect the watermark.

Tamper-resistance As well as requiring the watermark to be robust to legitimate signal distortions, a watermark may also be subject to signal processing that is solely intended to remove the watermark. In addition, when many copies of the same content exist with different watermarks, as would be the case when a watermark is used for buyer identification, further attacks are possible based on collusion amongst several buyers.

It is important that a watermark be resistant to tampering. There are a number of possible ways this may be achieved:

1. **Private watermark**: We believe that a private watermark, i.e. where either the decoder requires knowledge of the unwatermarked content or the pseudo-random noise sequence that constitutes the watermark is only known to the sender and receiver, are inherently more tamper resistant that public watermarks in which anybody is free to decode the watermark.

For the case in which only a single watermarked copy of the content is available, the only attack appears to be to add noise to the image in the hope of destroying the watermark. However, it can be shown that the magnitude of noise that needs to be added to be confident that the watermark is destroyed is so large that the image fidelity will be severely degraded. For the case of multiple watermarked copies of the same content, more powerful collusion attacks are possible, the most obvious being to average together all n copies.

In the case where all knowledge to decode the watermark is public, the most obvious attack is to simply invert the encoding process.

DOCKE.

- 2. Asymmetric encoder/decoder: If removal of a public watermark requires inverting the encoding process, then it is highly desirable to make the encoder as complex as possible, especially if the watermark is only to be applied once. However, if decoders must run in real-time, then it is necessary for the decoding process to be significantly simpler than the encoding.
- Bit rate The bit rate of a watermark refers to the amount of information a watermark can encode in a signal. This is especially important for public watermarks.
- Modification and multiple watermarks In some circumstances, it is desirable to alter the watermark after insertion. For example, in the case of digital video discs, a disc may be watermarked to allow only a single copy. Once this copy has been made, it is then necessary to alter the watermark on the original disc to prohibit further copies. Changing a watermark can be accomplished by either (i) removing the first watermark and then adding a new one or (ii) inserting a second watermark such that both are readable, but one overrides the other. The first alternative does not allow a watermark to be tamper resistant since it implies that a watermark is easily removable. Allowing multiple watermarks to co-exist is preferable and also facilitates the tracking of content from manufacturing to distribution to eventual sales, since each point in the distribution chain can insert their own unique watermark.
- Scalability In commercial applications, the computational costs of the encoder and decoder are important. In some applications, the insertion is only done once and can be performed off-line. Consequently, the cost of encoding may be less important than the cost of decoding, which may have to occur at real-time video rates, for example. Computational requirements constrain a watermark to be simple, but this simplicity may significantly reduce the resistance to tampering. Further, it is well known that computer speeds are approximately doubling every eighteen months, so that what looks computationally unreasonable today may very quickly become a reality. It is therefore very desirable to design a watermark whose decoder is scalable with each generation of computers. Thus, for example, the first generation of decoder might be computationally inexpensive but might not be as reliable as next generation decoders that can afford to expend more computation to deal with issues such as geometric distortions.

In the next section, we summarize early work on watermarking and then describe more recent work which attempts to insert a watermark into the perceptually significant regions of an image.

3. A FRAMEWORK FOR WATERMARKING

The process of watermarking an image can be represented by the addition of a noise term that is a function of the watermark signal, w, and possibly of the original image, I. The watermarked image, I' is then given by:

$$I' = I + f(I, w) \tag{1}$$

The watermarked image may then be subject to any number of distortions due to tampering or common use which can also be represented as a noise process, n. In many cases, the noise may be approximated by a linear additive process. However, distortions such as geometric transforms of an image may be highly non-linear and image dependent, i.e. n = n(I). The image presented at the decoder, I'', is then given by:

$$I'' = I' + n = I + f(I, w) + n(I)$$
⁽²⁾

At the decoder, we wish to extract the watermark signal, w, i.e. the unwanted signal (or noise) is the image, I. It should be noted that the magnitude of I is very much larger than the inserted watermark, F(I, w), and the distortions, n, otherwise the image fidelity would not be preserved. Consequently, the signal-to-noise ratio at the input to the decoder, where the signal is now the watermark, w, is much less than one. It is immediately apparent that methods that use the original image as part of the decoding process can greatly improve the SNR by simply subtracting the original image, I from (2).

There are many ways to characterize the numerous proposed watermarking methods. Two properties which we think are important are (1) whether the watermark is inserted into perceptually significant regions of the image and (2) whether the inserted signal, f(I, w), is independent of the image, I.

DOCKE







Figure 2. Block diagram of non-linear insertion method when the watermark signal is a function of the image.

Consider first the case in which the inserted signal is independent of the image, i.e. f(I, w) = w. In this case, Equation 2 reduces to:

$$I'' = w + I + n \tag{3}$$

where the signal is w and the noise is I + n, and the signal can be extracted using traditional matched filtering. In this case, if the watermark is to be placed in perceptually significant regions of the image, then it must be bandpass filtered based on exiting knowledge of the human auditory or visual systems. This is illustrated in Figure 1.

However, there is a (possible) disadvantage to shaping the watermark spectrum independently from the image to match currently known human auditory or visual systems. The power present in these frequency bands varies greatly from image to image. Consequently, if simple linear addition of the shaped watermark and image occurs then the magnitude of the watermark must be very low to avoid worst case scenarios in which the image energy in a particular band is very low and artifacts are created because the watermark energy was too strong relative to the image. Conversely, if the image energy is very strong in a particular channel, there is an opportunity to add relative more watermark energy without affecting the image fidelity.

Inserting a signal that is a function of the image leads to a non-linear insertion procedure, as illustrated in Figure 2. For example, Cox *et al*^{1,2} proposed scaling the watermark to a fixed fraction of the energy present in a particular frequency coefficient, such that:

$$I' = I(1 + \alpha w) \tag{4}$$

Such a procedure has the advantage that when the image energy in a particular frequency channel is small, the watermark energy is also reduced, thereby avoiding artifacts, and when the image energy is large, the watermark energy is increased, thereby improving the robustness of the procedure. In general, if the watermark is chosen so that its spectrum is white, then multiplication or scaling by the corresponding image coefficient can be thought of as shaping the watermark to the spectrum of the image. The α term then scales the shaped watermark to an acceptable level that is a compromise between robustness and perceptibility. Of course, shaping the watermark across the entire image spectrum, including perceptually insignificant regions, is unnecessary, and the two procedures should

DOCKET



Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts



Keep your litigation team up-to-date with **real-time** alerts and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research



With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips



Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.

