ELSEVIER

Signal Processing 81 (2001) 1251–1260

SIGNAL PROCESSING

www.elsevier.nl/locate/sigpro

Optimal transform domain watermark embedding via linear programming

S. Pereira*, S. Voloshynoskiy, T. Pun

University of Geneva-CUI, 24 rue General Dufour, CH 1211, Geneva 4, Switzerland

Received 15 April 2000; received in revised form 31 October 2000

Abstract

Invisible digital watermarks have been proposed as a method for discouraging illicit copying and distribution of copyright material. In recent years, it has been recognized that embedding information in a transform domain leads to more robust watermarks. A major difficulty in watermarking in a transform domain lies in the fact that constraints on the allowable distortion at any pixel may be specified in the spatial domain. The central contribution of the paper is the proposal of an approach which takes into account spatial domain constraints in an optimal fashion. The main idea is to structure the watermark embedding as a linear programming problem in which we wish to maximize the strength of the watermark subject to a set of linear constraints on the pixel distortions as determined by a masking function. We consider the special cases of embedding in the DCT domain and wavelet domain using the Haar wavelet and Daubechies 4-tap filter in conjunction with a masking function based on a non-stationary Gaussian model, but the algorithm is applicable to any combination of transform and masking functions. Our results indicate that the proposed approach performs well against lossy compression such as JPEG and other types of filtering which do not change the geometry of the image.

Keywords: Watermark; Linear programming; Copyright; DCT; Wavelet

1. Introduction

The idea of using a robust digital watermark to detect and trace copyright violations has stimulated significant interest among artists and publishers in recent years. Podilchuk [17] gives three important requirements for an effective watermarking scheme: transparency, robustness and capacity. Transparency refers to the fact that we would like the watermark to be invisible. The watermark should also be robust against a variety of possible attacks by pirates. These include robustness against compression such as JPEG, scaling and aspect ratio changes, rotation, cropping, row and column removal, addition of noise, filtering, cryptographic and statistical attacks, as well as insertion of other watermarks [15]. The other requirement is that the watermark be able to carry a certain amount of information i.e. capacity. In order to attach a unique identifier to each buyer of an image, a typical watermark should be able to carry at least 60–100 bits of information. However, most of the

* Corresponding author.

E-mail addresses: shelby.pereira@cui.unige.ch (S. Pereira), svolos@cui.unige.ch (S. Voloshynoskiy), thierry.pun@cui.unige.ch (T. Pun).

0165-1684/01/$ - see front matter
PII: S0165-1684(01)00042-1

work in watermarking has involved a one bit watermark. That is, at detection a binary decision is made as to the presence of the watermark most often using hypothesis testing [23]. Barni [1] encodes roughly 10 bits by embedding 1 watermark from a set of 1000 into the DCT domain. The recovered watermark is the one which yields the best detector response.

Watermarking methods can be divided into two broad categories: spatial domain methods such as [4,16] and transform domain methods. Transform domain methods have for the most part focused on DCT [17,1,14], DFT [13,2] and most recently wavelet domain methods [17,3,24]. Transform domain methods have several advantages over spatial domain methods. Firstly, it has been observed that in order for watermarks to be robust, they must be inserted into the perceptually significant parts of an image. For images these are the lower frequencies which can be marked directly if a transform domain approach is adopted [6]. Secondly, since compression algorithms operate in the frequency domain (for example DCT for JPEG and wavelet for EZW) it is possible to optimize methods against compression algorithms as will be seen in Section 3. Thirdly, certain transforms are intrinsically robust to certain transformations. For example, the DFT domain has been successfully adopted in algorithms which attempt to recover watermarks from images which have undergone affine transformations [13].

While transform domain watermarking clearly offers benefits, in some cases it is desirable to specify constraints in another domain (spatial or another transform domain). In this case, the problem is more challenging since it is more difficult to generate watermarks in one domain while taking into account constraints in another. For example, the problem arises since constraints on the acceptable level of distortion for a given pixel may be specified in the spatial domain. In the bulk of the literature on adaptive transform domain watermarks, a watermark is generated in the transform domain and then the inverse transform is applied to generate the spatial domain counterpart. The watermark is then modulated as a function of a spatial domain mask in order to render it invisible. However, this spatial domain modulation is suboptimal since it changes the original frequency domain watermark. In the case of a DFT domain watermark, multiplication by a mask in the spatial domain corresponds to convolution of the magnitude of the spectrum. Unfortunately, to correctly account for the effects of the mask at decoding a deconvolution problem would have to be solved. This is known to be difficult and to our knowledge in the context of watermarking this problem has not been addressed. Methods proposed in the literature simply ignore the effects of the mask at decoding. One alternative which has recently appeared is the attempt at specifying the mask directly in the transform domain and ignoring spatial domain masking [17]. However other authors (e.g. Swanson [20]) have noted the importance of masking in the spatial domain even after a frequency domain mask has been applied. It should be noted that masking in one domain is not easily formulated since defining a spatial mask influences a frequency mask and vice-versa.

In this publication, we develop a new approach for the mathematical modelling of the embedding process. In particular, we derive an optimized strategy for embedding a watermark in the wavelet and DCT domains when the masking constraints are specified in the spatial domain. In fact, the key idea is to optimize the encoding of the watermark with respect to the detector while using all available information about the image. This framework overcomes the problems with many proposed algorithms which adopt a suboptimal spatial domain truncation and modulation as determined by masking constraints. Furthermore, we will develop an algorithm which is image dependent. Unlike many of the embedding strategies described in the literature which treat the image as noise possibly modelled by a probability distribution, the algorithm we describe uses information about the image at embedding. We consider only the problem of generating watermarks which are robust against attacks that do not change the geometry of the image. We will work with an 80 bits watermark which corresponds to a capacity sufficient for most watermarking applications. We begin in Section 2 by presenting the spatial domain masking methods we adopt in the rest of the paper. In Section 3, the embedding algorithm is described and applied to the case of DCT domain embedding.