Digital Image Watermarking in the Wavelet Transform Domain

Diplomarbeit

zur Erlangung des Diplomgrades
an der Naturwissenschaftlichen Fakultät
der Universität Salzburg

eingereicht von

Peter Meerwald

Salzburg, am 11. Jänner 2001

Thanks

First of all, I would like to thank my parents for their care, support and advice.

It is a great pleasure to thank my advisor, Andreas Uhl, who has always been both, critical and supportive.

The discussion with Patrick Loo and the ongoing support has always been enjoyable and helpful. I want to thank all my colleagues and the local technical staff for their patience. Without the help and suggestions of Gabriele Prohaska, this work would be less organized, harder to understand and contain more errors, definitely.

Furthermore, I would like to thank the principal authors of and all contributors to The Gimp[1], LyX[2], Linux[3] and many other open-source projects – software which greatly facilitated the task of writing this thesis.

Part of work for this diploma thesis was funded by Prof. P. Zinterhof's research projects on "Object-based Image and Video Compression with Wavelets" (FWF grant P13732) and "Adaptive, Hybrid and Parallel Wavelet Image Coding" (FWF grant P11045).

[1] The GNU Image Manipulation Program, available at `http://www.gimp.org`.
[2] The LaTeX Document Processor, `http://www.lyx.org`.
[3] The latest Linux kernel is always at `http://www.kernel.org`.

Contents

List of Figures

List of Tables

List of Algorithms

Abstract

The development of compression technology – such as the JPEG, MPEG and more recently the JPEG2000 [1] image coding standards – allowed the widespread use of multimedia applications. Nowadays, digital documents can be distributed via the World Wide Web to a large number of people in a cost-efficient way. The increasing importance of digital media, however, brings also new challenges as it is now straightforward to duplicate and even manipulate multimedia content. There is a strong need for security services in order to keep the distribution of digital multimedia work both profitable for the document owner and reliable for the customer. Watermarking technology plays an important role in securing the business as it allows to place an imperceptible mark in the multimedia data to identify the legitimate owner, track authorized users via fingerprinting [52] or detect malicious tampering of the document [120].

Previous research [39] indicates that significant portions of the host image, e.g. the low-frequency components, have to be modified in order to embed the information in a reliable and robust way. This led to the development of watermarking schemes embedding in the frequency domain. Nevertheless, robust watermarking in the spatial domain can be achieved [20] at the cost of explicitly modeling the local image characteristics. These characteristics can be more easily obtained in a transform domain, however.

Many image transforms have been considered, most prominent among them is the discrete cosine transform (DCT) which has also been favored in the early image and video coding standards. Hence, there is a large number of watermarking algorithms that use either a block-based [110, 18] or global DCT [39, 6]. Other transforms that have been proposed for watermarking purposes include the discrete Fourier transform (DFT) [199], the Fourier-Mellin transform [175] and the fractal transform [193, 57]. In this work, we focus on the wavelet domain for the reasons given below.

With the standardization process of JPEG2000 and the shift from DCT- to wavelet-based image compression methods, watermarking schemes operating in the wavelet transform domain have become even more interesting. New requirements such as progressive and low bit-rate transmission, resolution and quality scalability, error resilience and region-of-interest (ROI) coding have demanded more efficient and versatile image coding [27]. These requirements have been met by the wavelet-based "Embedded Block Coding with Optimized Truncation" (EBCOT) system [232], which was accepted with minor modifications as the upcoming JPEG2000 image coding standard [1]. The wavelet transform [47] has a number of advantages [266, 151] over other transforms such as the DCT that can be exploited for both, image compression and watermarking applications. Therefore, we think it is imperative to consider the wavelet transform domain for watermarking applications.

In chapter 1, we will try to motivate the watermarking research effort and discuss some of the applications that require watermarking technology. The "watermarking problem" is introduced in chapter 2 and we state why it is hard to design watermarking algorithms that can fulfill the requirements derived in the previous chapter. Furthermore, we investigate the relationship of watermarking to cryptography and image compression and justify the choice of the wavelet transform. Hereby, the required background in wavelet theory and perceptual coding techniques is briefly introduced.

Chapter 3 starts by first characterizing the most important and distinguishing features of previously proposed wavelet-based watermarking schemes. We organize the overwhelming amount of algorithms proposed in the literature in two main categories: additive and quantize-and-replace strategy embedding techniques. Further on, each approach is discussed in detail, building on the experience that was gained from implementing some of the watermarking schemes.

Our own contributions are presented in chapter 4. First, we propose using wavelet filter parametrization as a means to improve the security of many previously proposed algorithms and demonstrate that our concept of secret key-dependent wavelet filters can be easily employed as a security framework. Motivated by the increasing importance of the forthcoming JPEG2000 image standard, we present a novel watermarking schemes that is compatible with the independent code-block processing approach of the new image coding standard. Two application scenarios, copyright protection and image authentication, are considered and we demonstrate that our quantization-based embedding technique can successfully encode and decode a binary watermark on-the-fly in the JPEG2000 coding process.

In chapter 5, a classification of attacks on watermarks is given which will be used to discuss the robustness results in chapter 6.

Parts of this thesis have been presented at the following conferences and published in the conference proceedings:

- Peter Meerwald and Andreas Uhl. Sicherheit und Robustheit Wavelet-basierter Watermarking-Algorithmen. In M. Schumacher and R. Steinmetz, editors, *Tagungsband des Workshops Sicherheit in Netzen und Medienströmen*, Springer-Verlag, Informatik aktuell, pages 181–190, Berlin, Germany, September 2000.

- Peter Meerwald and Andreas Uhl. A survey of wavelet-domain watermarking algorithms. In *Proceedings of SPIE, Security and Watermarking of Multimedia Contents III*, volume 4314, San Jose, CA, USA, January 2001.

The following papers have been submitted to conferences:

- Peter Meerwald. Quantization watermarking in the JPEG2000 coding pipeline. *submitted to CMS '01*, Darmstadt, Germany, May 2001.

- Peter Meerwald and Andreas Uhl. Watermark security via wavelet filter parametrization. In *submitted to ICIP '01*, Thessaloniki, Greece, October 2001.

Chapter 1

Introduction

"Information is inevitably physical"
Rolf W. Landauer (1927 - 1999)

1.1 Intellectual Property and the Digital Age

With the growth of the Internet and the immediate availability of computing resources to everyone, “digitized property” can be reproduced and instantaneously distributed without quality loss at basically no cost. Until now, intellectual property (IP) and value has always been bound to some physical container that could not be easily duplicated, thereby guaranteeing that the creator benefits from his work.

Barlow [3] and Dyson [58] consider traditional copyright law inappropriate for the “digital age” and suggest to overcome the restrictions and problems by associating value not to digital content itself but mainly to service and personal 'experience' built around to it. Since personal experience can hardly be duplicated over the Internet, there is no need for extensive regulations.

Clearly, there are businesses like the music or photography industry that can not adopt this paradigm since they trade basic content and therefore have to stick with traditional copyright enforcement to guarantee income. As audio, video and other works become available in digital form, it may be that the ease with which perfect copies can be made will lead to large-scale unauthorized copying which will undermine the music, film, book and software publishing industries.

One technical way to make law enforcement and copyright protection for digital media possible and practical is digital watermarking which is aimed to automatically detect and possibly also prosecute copyright infringement. There has therefore been significant recent research into “watermarking” (hiding copyright messages) and “fingerprinting” (hiding serial numbers or a set of characteristics that tend to distinguish an object from other similar objects); the idea is that the latter can be used to detect copyright violators and the former to prosecute them.

Year	1992	1993	1994	1995	1996	1997	1998	1999	2000
Publications	2	2	4	13	29	64	103	200+	150+

Table 1.1: Number of publications in the watermarking field during the past years, according to [57] and our own research bibliography.

Watermarking is a relatively young research field. In spite of the very active[1] research (see table 1.1) and the heavy industrial demand, successful real-world applications have not been developed yet. Petitcolas [179, 178] has shown that commercial image watermarking applications available today can be easily attacked.

The music industry is about to set a new standard for compressed audio files in order to replace the ubiquitous but unprotected MPEG audio streams[2]. Likely, audio watermarking will become the key technology in that effort. Similarly, DVD technology depends on video watermarking for copy protection and copy management [95, 160, 14, 139].

1.1.1 Copyright Protection

The goal of watermarking for copyright protection is to embed a "mark" into the image data that can identify the copyright holder of the work. Together with owner identification, one might also want to embed a mark (or fingerprint) identifying the buyer of a work for circulation tracking. The mark can be a registered number (like the UPC[3] found on compact disk media), a text message or graphical logo, or some unique pattern (similar to a DNA fingerprint). The term watermark stems from the ancient art of marking paper with a logo for the same purpose.

Digital watermarks can either be perceptible or imperceptible. Visible image watermarks, often the logo of the copyright holder, can be easily applied to the image but are hard to remove. Mintzer describes an successful implementation of visible image watermarking in [162, 163]. Many applications require the watermark to be invisible, however. This work focuses on invisible watermarks in digital images only.

The embedded, invisible watermark has to be robust against common image processing operations like image compression (e.g. JPEG), image filtering (edge enhancement, contract enhancement, ...), and geometrical transformations (e.g. cropping, scaling, ...). Therefore, the watermark can not be stored in the file format, but has to be embedded into the image data itself. In order to establish a proof of ownership in a trial, a watermarking scheme also has to be secure against intentional malicious attacks; here, cryptographic techniques and statistical properties of pseudo-random numbers play an essential role.

[1] In 1998 more than one-hundred papers have been published, see also [57]. The trend continued 1999 and 2000 with more than 200 and 150 papers published, respectively.

[2] MPEG-1 audio layer 3, commonly called MP3

[3] UPC ... Universal Producer Code

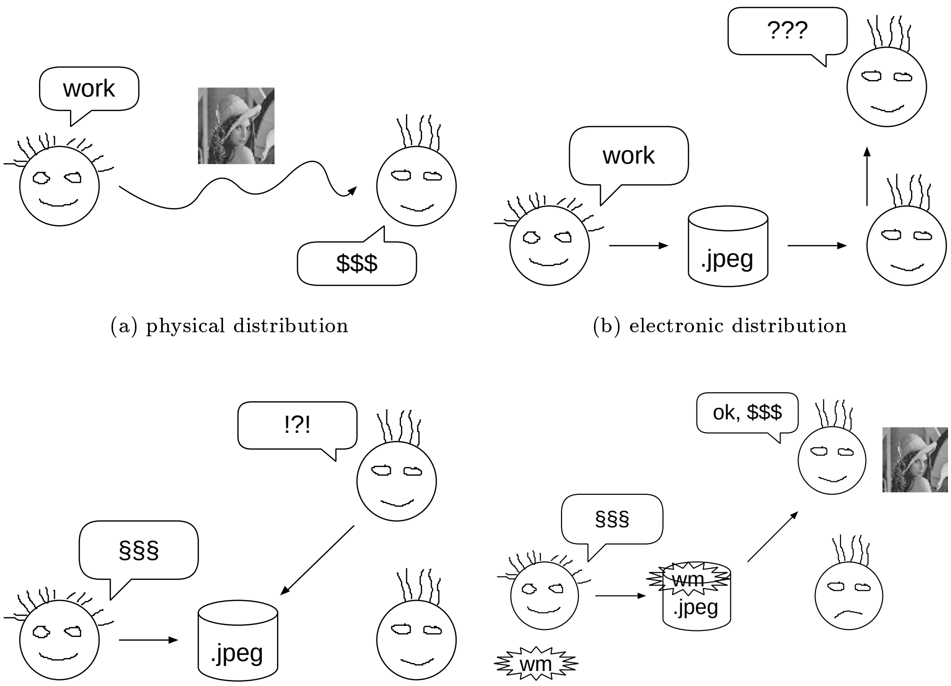

(a) physical distribution

(b) electronic distribution

(c) legal confrontation

(d) protection by watermarking

Figure 1.1: Intellectual property transfer between the creator and the customers of a work. (a) The photography is distributed traditionally and therefore hard to duplicate or manipulate. The clients pay for the work of the IP creator. (b) The work is distributed in electronic, digitized form. Making copies is cheap and easy, there seems to be no reason for a third party to pay royalties to the creator of the work. (c) The creator has difficulties to track copies of his digitized work and claim ownership in a legal trial. (d) A watermark can be used to convince the IP customers to pay the royalties without limiting usage of the work. In addition, a watermark may provide extra information and guarantee data integrity.

1.1.2 Image Authentication and Data Integrity

Another application of watermarking is image authentication and "tamper detection". Digital photographs are being used more and more often as court evidence nowadays. Here, watermarking is used to detect significant modification of the image. Digital images are susceptible to seamless modifications from sophisticated image processing applications. Watermarks can be used here as a means to verify the genuineness of an image. Verification watermarks are required to be fragile, so that any modification to the image will destroy (or detectable alter) the mark. Unlike cryptographic message digests which can only validate identical copies, watermarking for image authentication should tolerate some well-defined image distortion (e.g. file format conversion, re-sampling, re-compression or progressive transmission).

1.1.3 Data Hiding and Image Labeling

Data hiding or steganography tries to invisibly embed the maximum amount of data into a host signal (e.g. an image). This allows communication using often enciphered messages without attracting the attention of a third party. Typically, robustness requirements are low for steganographic purposes, instead invisibility and capacity are of prime importance.

Image labeling is an application where information about the image content is encoded as a watermark and inserted into the image to assist image retrieval from a database or provide extra information to the viewer.

1.1.4 Watermarking Everything

Watermarking, that is the technique of placing and transmitting small amounts of data imperceptibly in host or cover data, has recently found many new applications. However, steganography and data hiding has been studied long before [209, 236] and the use of paper watermarks for copy protection can be traced back until the 15th century[4] (see figure 1.3 (a)). Surprisingly, ancient works first prompted for a technical solution for copyright protection of digital images as soon as they were displayed in digital libraries[5] available through the Internet (see figure 1.3 (b)).

Nowadays, there exist watermarking methods for virtually every kind of digital media: text documents [157, 141, 19, 13], images (most research concentrates on this media type, see Nikolaidis [168] overview of image security techniques), video [75, 129, 194, 96, 138, 228], audio [17, 9, 192, 165, 229, 233, 2, 225, 148, 133, 103], even for 3D polygonal models [11, 99, 213, 169], maps [101] and computer programs [32, 33]. Interestingly, watermarking technology is not limited to digital media , but also applicable to e.g. chemical data like protein structure [82, 61].

---

[4] See the Gravell watermark archive, available at `http://ada.cath.vt.edu:591/DBs/Gravell/default.html`.

[5] The IBM Digital Library project, see `http://www.dlib.org/dlib/july97/vatican/07gladney.html`.

(a) Clinton and Monica

(b) Clinton and Hillary

Figure 1.2: Data integrity: hard to judge from a digital photography with the naked eye, is (a) or (b) a trustworthy photography?
(Images by Ching-Yung Lin, see http://www.ctr.columbia.edu/~cylin/auth/auth.html.)

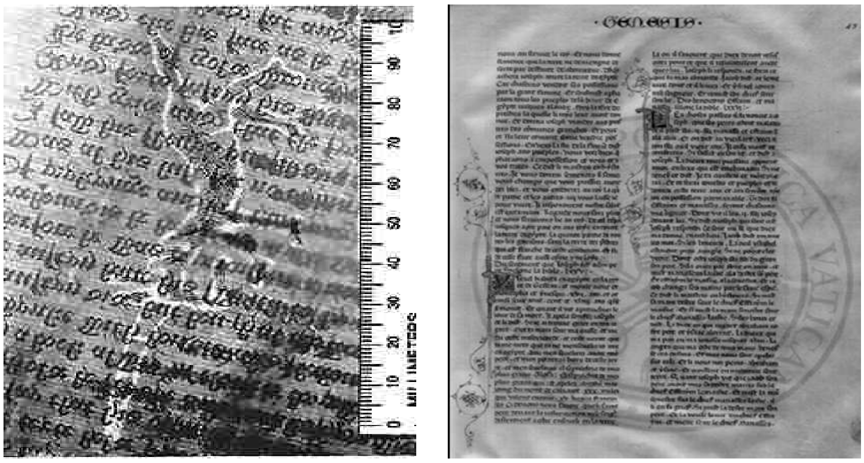

(a) historic watermark (b) modern watermark

Figure 1.3: An historic unicorn watermark (a) in William Caxton's edition of the Canterbury Tales, 1477. (b) A visible watermark in a French 15th century Genesis excerpt, created by the IBM Vatican Library project.

In this work we concentrate on digital image watermarking. Most of the algorithms examined in this work operate in the wavelet transform domain which seems to be an excellent choice for image processing and image compression in general and image watermarking in particular. See section 2.5.3 and section 2.7 for the rational of this choice.

Since the upcoming image compression standard, JPEG2000, will be based on the wavelet transform, as opposed to the DCT the JPEG standard is built on, it seems only natural to take advantage of the superior performance and modeling properties of the wavelet transform for watermarking purposes as well.

Chapter 2

The Watermarking Problem

Image watermarking imperceptibly embeds data into a host image. The general process of image watermarking is depicted in figure 2.1.The original image (host image) is modified using the signature data to create the watermarked image. In this process, some error or distortion is introduced. To ensure transparency of the embedded data, the amount of image distortion due to the watermark embedding process has to be small. The watermarked image is then distributed and may circulate from legitimate to illegitimate customers. Thereby, it is subjected to various kinds of image distortion. Image distortion may result from e.g. lossy image compression, re-sampling or from specific attacks on the embedded data.

Note, that we do not discuss visible watermarking in this work. The methodology for visible watermarking is very different from invisible watermarking. For visible watermarking techniques, see for example Mintzer's and Gladney's [73, 74] description of the IBM digital library project [162].Our focus is on invisible, or better, imperceptible watermarks.

The extraction process may or may not, depending on the nature of the application, require knowledge of the original host image to estimate the hidden signature from the distorted image that is received. The watermark is recovered from the host image. It is desired that the difference between the extracted and the original signature is as low as possible.

2.1 The Watermarking process

In order to see the different aspects of the watermarking problem, depending on the particular applications and the applications' requirements, we have to refine the general watermarking model (figure 2.1) and have a closer look at the successive stages of the watermarking process. These stages comprise

- the embedding stage (figure 2.2),
- the extraction stage (figure 2.4),

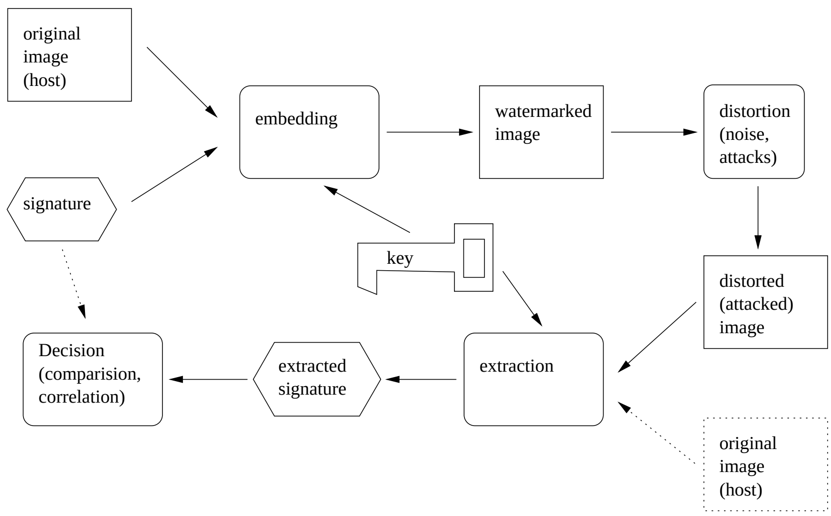

Figure 2.1: The data hiding model, a general overview.

- the distribution stage (figure 2.3) and
- the decision stage (figure 2.5).

2.1.1 Embedding stage

Except for some very early watermarking schemes such as the Patchwork approach [10], Walton's image authentication method [249] or Pitas's bi-polar watermarking scheme [183], all robust watermarking algorithms operate in a transform domain that offers access to the frequency components of the host image. By omitting the transform step and performing data embedding step in the spatial domain, one can design a simple and computational efficient algorithm for watermarking. However, these approaches failed to achieve good robustness and sacrifice strength against compression attacks.

In the embedding stage, the host image is therefore first transformed to a domain that facilitates data embedding. This work exclusively considers the wavelet and wavelet packet transform domains. Other commonly used frequency domain representations can be obtained by the DCT (discrete cosine transform) or the DFT (discrete Fourier transform). Section 2.5 discusses some of the rational that makes an image's frequency representation a favorable playground for watermark embedding.

The signature data (also called the message) can be some binary data, a small image (a "logo") or a seed value to a pseudo-random number generator to produce a sequence of numbers with a certain distribution (e.g. Gaussian or uniform). Typically, the signature data has to be encrypted to decorrelate the information and subjected to some error-correcting coding scheme.

Next, the subset of the transform coefficients is modified with the prepared signature data. Optionally, we can employ a model of human perception to

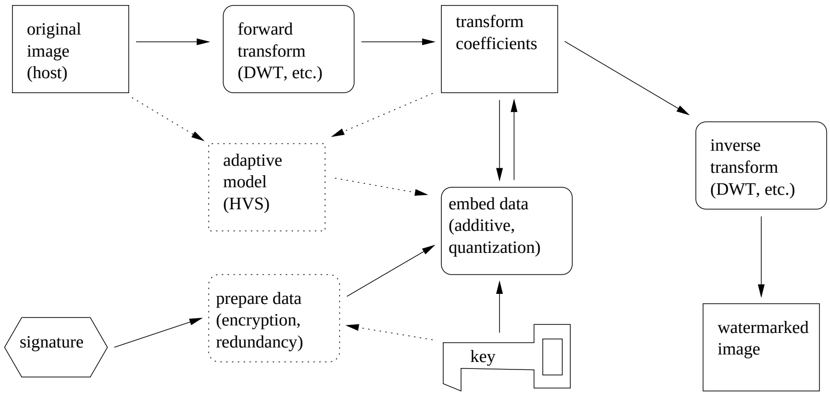

Figure 2.2: Model of the watermark embedding stage.

weight the strength of the embedding modifications. Note that by choosing a suitable frequency transform domain and selecting only certain coefficients (typically in low- to mid-frequency range, see for example the work of Cox [39]), a lot of the human visual system (HVS) modeling can be done implicitly. The better the image transform approximates the properties of HVS, the easier it is to put more energy in the embedded signal without causing perceptible distortion. See section 2.6 for more details about modeling certain properties of the human visual system.

Finally, the inverse transformation is applied on the modified transform domain coefficients to produce the watermarked image.

2.1.2 Distribution

The watermarked image is then distributed – for example published on a web server or sold to a customer. Nowadays, distribution of digital media often includes lossy compression prior to transmission. The impact of compression on the embedded watermark data is discussed in section 2.7.3.

During transmission and distribution of the watermarked image, not only compression adds distortion to the host data, but also transmission errors and common image processing tasks, such as contrast enhancement, re-sampling and gamma correction, contribute errors to the watermarked image. Especially geometric image manipulation like scaling, rotation or cropping has been proved to be very harmful to the embedded watermark. All manipulation of the watermarked image data has to be seen as an attack on the embedded information. Modifications that occur during normal image processing are called coincidental attack, while attacks that attempt to weaken, remove or alter the watermark itself are termed hostile or intentional attacks. In chapter 5 we characterize a number of attacks and describe counter-measures that can be taken by the watermarking system.

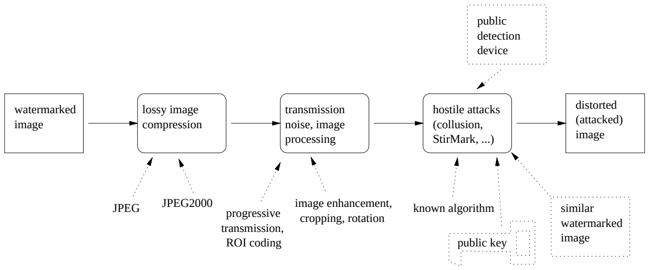

Figure 2.3: Model of the distribution of the watermarked image.

2.1.3 Extraction stage

Eventually, after the watermarked image has undergone severe distortion as described in the previous section, one would like to extract the embedded signature from the host data. This can be done by the party that embedded the watermark, the customer that received the image, a designated party – such as a Web crawler that scan the Internet for illegal copies of the protected work [272] or a legal prosecution official – or by a third party. In the first case, the secret key used to embed the watermark as well as the original image might be available. This tremendously facilitates the watermarking system and makes watermark detection relatively straightforward. We call detection systems that have access to the secret (private) key and original image non-oblivious, non-blind or private watermarking systems.

The other extreme is the case where neither the private key nor the original image is available during the extraction process. These watermarking systems are called public key watermarking systems. However, no reliable public-key watermarking system is known to work and it is likely that no such system can ever be built [98]. Recently, also asymmetric watermarking schemes have been proposed that use different keys for embedding and detecting the watermark [62]. The relationship between watermarking and public-key cryptography is explained in section 2.4.

A watermarking scheme that allows to extract the signature data without reference to the original, unwatermarked, image host is dubbed blind or oblivious watermarking scheme. There are also detection or extraction methods that rely on some data or features derived from the original host image. These schemes have been named semi-blind or semi-oblivious watermarking algorithms.

To summarize these important distinctions based on the availability of the original image, there are

- blind or oblivious,
- semi-blind and
- non-blind or non-oblivious

watermark extraction methods. Regarding the key material necessary for watermark extraction we can distinguish

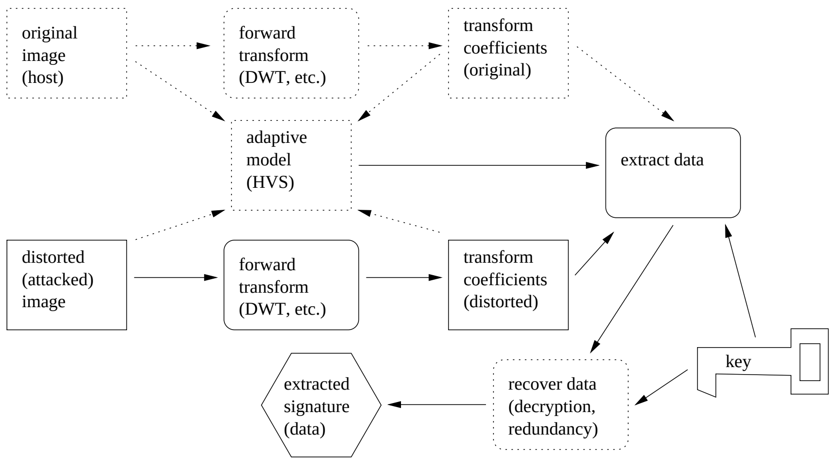

Figure 2.4: Model of the watermark extraction stage.

- private key
- public key
- and asymmetric

watermarking schemes.

2.1.4 Decision stage

In the decision stage, the watermarking system analyzes the extracted data. Depending on the type of the application and the nature of the signature data, the decision stage can produce a number of different outputs.

For image copy protection applications, the output of the watermarking system can range from simple to more complicated answers. In the simplest case, the result is just a yes/no decision indicating if the copyright holder's mark has been found in the received image data. More complex systems return the embedded logo image or the textual copyright information that was placed into the host image data. A widely used similarity measure between the original, W, and the extracted watermark sequence, $W^\star$, is the normalized correlation for pseudo-random sequences,

$$\delta = \frac{W^\star \cdot W}{\|W^\star\| \cdot \|W\|},$$

or the Hamming distance for binary messages, $w_i \in \{-1, 1\}$,

$$\delta = N - \sum^{N} w_i^\star \cdot w_i.$$

The extracted watermark yes/no answer can be derived from the similarity measure δ with an appropriate threshold τ, i.e. if $\delta \geq \tau$ then is watermark is detected otherwise watermark could not be found in the image.

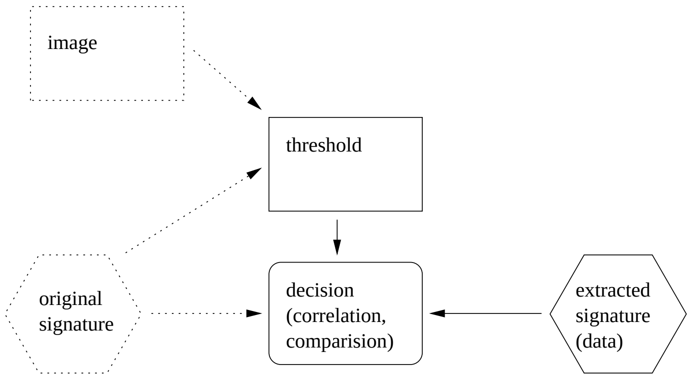

Figure 2.5: Model of the decision stage.

Image labeling and data hiding applications will typically try return the message originally embedded. Since message corruption can not be tolerated, the use of error-correcting codes is mandatory for this type of application.

Watermarking schemes for image authentication and data integrity verification will either just return a yes/no answer to indicate if the image data has been tampered with, return the identify of the legitimate source or try to identify the image regions that have been adversely manipulated (e.g. telltale watermarking [118, 120]).

2.2 Application Aspects and Requirements

Different watermarking application have different requirements. In the following, we present some application scenarios described by Piva [186], Hartung [76] and other authors.

For image data authentication, the embedded watermark has to be invisible to a human observer and it should be altered (or broken) by virtually any intentional modification of the image. Furthermore, it should be difficult to insert a false watermark and the watermarking scheme should be able to indicate regions where alterations in the image have taken place.

Several image copyright protection application scenarios are possible. First, the owner of an image can embed an invisible, robust and quickly extractable watermark to identify unauthorized copies. Employing a web-crawler for this task has been controversially discussed in the literature [273].

Second, demonstration of ownership requires in addition to robustness that the watermarking scheme is also non-invertible, bind and private. These constraint have been carefully analyzed [43]. On the other hand, speed and ease of detection or extraction is not of great relevance.

Finally, the copyright holder (the seller of an image) might also want to know which customer leaked an unauthorized copy of the data. Here, fingerprinting and circulation tracking techniques come into play to identify not only the seller but also the buyer of an image. To this aim, some additional requirements are

Terms used in this work	Synonymous terms
watermark embedding	casting,engraving,etching
watermark extraction	recovery, detection
host image	cover image
signature	embedding message, watermark
watermarked image	stego image
blind watermark	oblivious, public watermark
non-blind watermarking	non-oblivious, private watermarking

Table 2.1: Watermarking terminology.

necessary. For instance, it should be possible to generate a large number of different watermarks and the insertion of multiple watermarks should be handled properly.

For data hiding and image labeling purposes, the maximum capacity of embedded message is of prime importance. While certain unobtrusive distortions for acceptable for e.g. copyright protection schemes, there are much higher imperceptibility requirements for steganographic applications. Image labeling techniques require highly localized embedding of watermark information (preferable image object-based), which rules out methods that operate on the entire image [124].

As for cryptography, watermarking methods have to obey the Kerckhoff principle [100] which means that security and robustness claims have to take into consideration that the algorithms for watermark embedding and extraction are known in detail. Despite this insight, most advanced watermarking schemes today do not disclose the implementational details to reproduce the results. Furthermore, very little research effort has been devoted to analyze the security of watermarking schemes.

The message capacity, that is the number of bits that can be reliably embedded in the image data, is fairly limited. For copyright protection applications that involve identification of the copyright holder as well as the identification of the licensee of the image, different lower capacity bounds have been proposed by Piva [186] (300 bits) and Kutter [126] (64 bits); the later proposal stems from the ISO multimedia license plate standard[1].

2.3 Terminology

The watermarking and data hiding problem has been examined by various research communities (such as image processing, communication and information theory, cryptography, ...), each from a slightly different point of view. No standard terminology has been coined yet, although most of the approaches so far seem to share a common core model [180].

[1]ISO document 10919-4.

2.4 Relationship to Cryptography

Image cryptography is considered as an encoding technique for data transmission through communication channels under condition that a third party could not read and interpret the data [80]. However, transmitted data, especially in scrambled form, can attract attention and impel law-enforcing authorities to take a closer look [41]. Nevertheless, cryptography has become one of the main tool for privacy, access control, authentication, digital signatures and secure messaging.

Steganography implants the secret message in some form of cover data, typically digital images or video streams, concealed as 'noise'. Without the correct key, it is virtually impossible to extract the hidden message or even detect its presence. This places an additional burden on the cryptanalyst who now has to examine unsuspicious-looking data for embedded steganographic messages [68].

Steganographic messages are usually encrypted in order to increase security, but also to conceal any statistical significant patterns. For image-type messages, mixing systems based on toral automorphisms [247, 248] or Kolmogorov flows [206] can be used.

The second relationship between cryptography and watermarking stems from the shared semantics of public- and private-key crypto-systems versus public- and private-key watermarking systems. In private-key crypto-systems, the same key is used to encrypt and decrypt a message (symmetric cipher), while in public-key crypto-systems, the keys for encryption and decryption are different (asymmetric cipher).

2.4.1 Public-key Watermarking

In a public-key watermarking system, a digital object is marked with the private key but the presence of a watermark can be checked using a public key. Of course, computation of the corresponding private key is infeasible, despite the availability of the public key and the algorithm of the watermarking system. The public key only allows to read the watermark, it can not be used to remove or forge a watermark. Public-key watermarking generally assumes that the unmarked original host is not required in the detection or extraction process (blind recovery).

Traditional watermarking systems using symmetric, private keys almost always allow to remove or insert forged watermarks. Public-key schemes have to permit secure watermark verification by third persons. However, no public-key watermarking schemes is known to exist, since most current approaches can not withstand public detector device attacks as described by Kalker and Linnartz [98, 140].

2.4.2 Asymmetric watermarking

Asymmetric watermarking techniques do not refer to the original image and employ different parameters than the ones used in the embedding process. The terms public-key watermarking and asymmetric watermarking are often used in the same context. If you take the terms literally you might differ between asymmetric watermarking, the keys enabling watermark detection differ from the keys needed for embedding the watermark, and public-key watermarking, the keys necessary for watermark detection are publicly known, thus, enabling everybody – including potential attackers – to detect an embedded watermark. Asymmetric watermarking might actually be a way of realizing a public-key watermarking system.

The public detector device attack described by Kalker and Linnartz [98, 140] which is based on the linearity of the detection process, does not work on asymmetric schemes, e.g. see Eggers [64, 62].

2.4.3 Visual Cryptography

Visual cryptography is a type of a cryptographic scheme to conceal images without any cryptographic computations. It is a visual variant of the k out of n secret sharing problem. One would produce transparencies that contain parts of the secret. Any k of the n transparencies stacked on a heap would reveal the secret, but less than k transparencies do not reveal any information. See figure 2.6[2] for an example for $k = 2$.

Contrary to steganography, there is no host data in visual cryptography. The secret is shared and can be extracted by combining part of the keys. The keys have visual representations (transparencies). See the papers by Naor [164], Droste [54] and Stinson [215].

2.5 Operating in the Transform domain

Transform domain watermarking techniques apply some invertible transform to the host image before embedding the watermark. Then, the transform domain coefficients are modified to embed the watermark and finally the inverse transform is applied to obtain the marked image. The transforms commonly used for watermarking purposes are the discrete cosine transform (DCT), the discrete Fourier transform (DFT), the fractal transform and the discrete wavelet transform (DWT) but there are also approaches dealing with more "exotic" transforms such as the Fresnel transform, the complex wavelet transform (CWT), the Fourier-Mellin transform and others.

Transform domain watermarking algorithms possess a number of desirable properties [30]: Since the watermark embedded in the transform domain is irregularly distributed over the area of local support after the inverse transformation, these methods make it more difficult for an attacker read or modify the mark. For watermarking strategies that depend on the global DCT this means the watermark is spread over the entire image. Of course, the wavelet transform or a block-based DCT only affects the local region. Furthermore, the frequency representation of the images allows to select only certain bands of the host signal

[2] From `http://cacr.math.uwaterloo.ca/~dstinson/visual.html`.

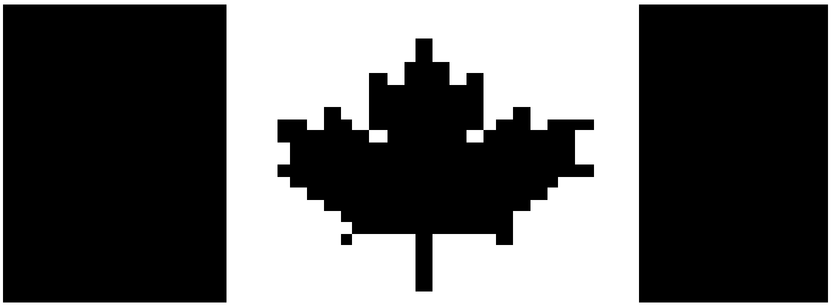

(a) original image

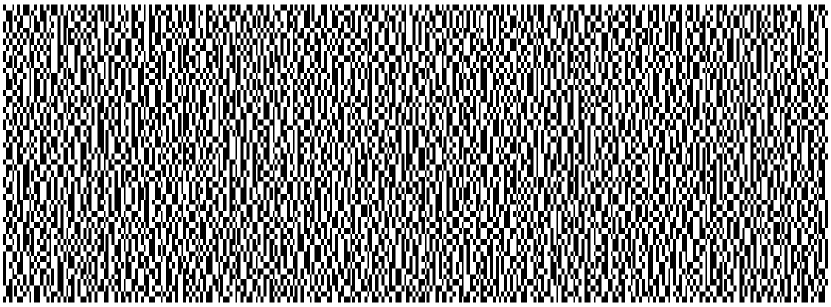

(b) secret share no. 1

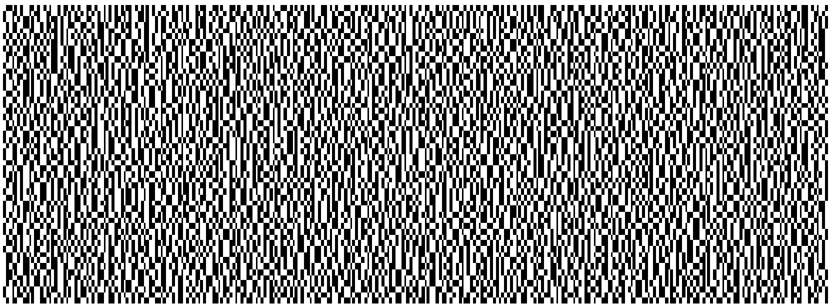

(c) secret share no. 2

(d) combining no. 1 and 2

Figure 2.6: Visual cryptography. The secret information (a) is split into two shares, (b) and (c). Only when combining the shares (d), the hidden information is revealed.

for watermarking. The human visual system has been observed to process certain frequency bands individually which led to the development of visual models that try to capture these characteristics [153, 257].

Algorithms operating in the frequency domain usually add the mark – or its spread spectrum signal – to a small subset of transform coefficients of the low or medium frequency range [39]. In spread spectrum techniques, a narrow-band signal which represents the message to be transmitted is modulated by a broadband carrier signal, which broadens or spreads the original, narrow-band spectrum; hence the term "spread spectrum".

2.5.1 Spread spectrum

The following properties of spread spectrum are particularly well-suited for watermarking [220, 77]:

Anti-jamming. The anti-jamming property results from the fact that an attacker does not know the privileged information that the sender and an authorized receiver possess. As a result, the attacker must jam the entire spectrum of the broadband signal. The jammer has limited power, however, so it can only jam each frequency with low power. Hence, the sender and receiver have an effective signal-to-jammer advantage (called the processing gain).
With application to watermarking, the anti-jamming property means that, in order to jam a watermark, an attacker must distort the marked media severely – so severely that the attacked media is no longer of acceptable quality or has no commercial value.

Low probability of intercept. The low probability of intercept property is a consequence of spreading: a large signal power is distributed over the entire frequency spectrum, so only a small amount of power is added at each frequency. Often, the increase is below the noise floor, so an attacker may not even detect the transmission of a spread-spectrum signal. This allows a watermark to be embedded unobtrusively.

Pseudo-noise. For security, the carrier is often a pseudo-noise signal, meaning that it has statistical properties similar to those of a truly random signal, but it can be exactly regenerated with knowledge of privileged information. For example, the carrier could be the output of a random-number generator that has been initialized with a particular seed, and the seed is known only to the owner.
The pseudo-noise property is useful for watermarking, because it makes it difficult for an attacker to estimate the watermark from marked media. In addition, with properly chosen pseudo-noise signals, even if the attacker can perfectly estimate some small segments of the watermark, it is not possible to determine the rest of the mark.

The low- and mid-frequency components of the image data represent most of the perceptual important information. Therefore, compression schemes and other image processing operations can hardly affect this significant portion of

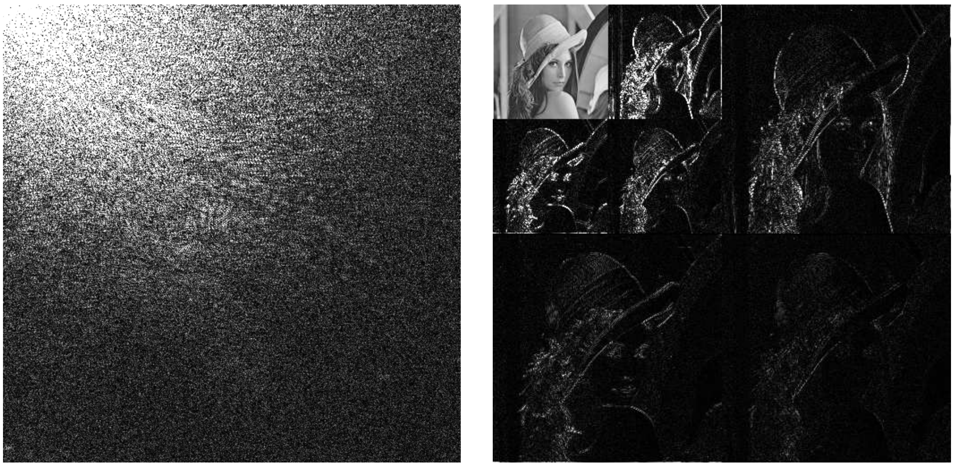

(a) DCT (b) DWT

Figure 2.7: Energy distribution in the transform domain. The "Lena" image has been transformed with the DCT (a) and a three-level DWT using bi-orthogonal 7/9 filters (b).

the host image without destroying much of the visual content of the image. Thus, adding the watermark in significant coefficients of the transform domain generally improves robustness. Spatial domain watermarking methods have to indirectly model the low-frequency component of the host image signal, which can be quite complicated to achieve [20].

In the DCT domain, the energy concentrates in the low frequency regions around the upper-left corner. The multi-resolution DWT representation has the low-frequency components of the image signal in the approximation subband, also located in the upper-left corner, while the high-frequency components are represented in the detail subbands at several resolutions (see figure 2.7). Most energy of the detail subbands is situated in edge areas and textured regions. Also note the similarity between subbands at different resolutions and orientations. The wavelet transform is presented in more detail in the next section.

The main disadvantages of frequency transform domain techniques are their computational cost, and, in the case of a global transform, their problem to adapt the watermark strength to the local image activity, making it more difficult to exploit certain characteristics of the HVS such as masking effects. The later shortcoming can be resolved by using the wavelet transform which provides both, frequency and spatial information of the host image.

2.5.2 The Wavelet Transform

The wavelet transform has been extensively studied in the last decade [46, 47]. Many applications of the wavelet transform, such as compression [132, 212, 202, 48], signal analysis and signal processing [108] have been found. There are many good tutorial books [260] and papers on this topic. Here, we just introduce the necessary concepts of the DWT for the purposes of this work.

The basic idea of the DWT for a one dimensional signal is the following. A signal is split into two parts, usually high frequencies and low frequencies. The edge components of the signal are largely confined in the high frequency part. The low frequency part is split again into two parts of high and low frequency. This process is continued until the signal has been entirely decomposed or stopped before by the application at hand. For compression and watermarking application, generally no more than five decomposition steps are computed. Furthermore, from the DWT coefficients, the original signal can be reconstructed. The reconstruction process is called the inverse DWT (IDWT).

Mathematically, the DWT and IDWT can be stated as follows. Let

$$H(\omega) = \sum_k h_k \cdot e^{-jk\omega},$$

and

$$G(\omega) = \sum_k g_k \cdot e^{-jk\omega}$$

be a low-pass and a high-pass filter, respectively, which satisfy certain conditions for reconstruction stated later. A signal, $F(n)$ can be decomposed recursively as

$$f_{j-1}^{low}(k) = \sum_n h_{n-2k} f_j(n)$$

and

$$f_{j-1}^{high}(k) = \sum_n g_{n-2k} f_j(n)$$

for $j = J+1, J, \ldots, J_0$ where $f_{J+1}(k) = F(f)$, $k \in Z$. $J+1$ is the highest resolution level index and J_0 is the low resolution level index. The coefficients

$$f_{J_0}^{low}(k), f_{J_0}^{high}(k), f_{J_0+1}^{high}(k), \ldots, f_J^{high}(k)$$

are called the DWT of the signal $F(n)$, where $f_{J_0}^{low}(k)$ is the lowest resolution part of $F(n)$ (the approximation) and the $f_j^{high}(k)$ are the details of $F(n)$ at various bands of frequencies. Furthermore, the signal $F(n)$ can be reconstructed from its DWT coefficients recursively,

$$f_j^{low}(n) = \sum_k h_{n-2k} \cdot f_{j-1}^{low}(k) + \sum_k g_{n-2k} \cdot f_{j-1}^{high}(k).$$

To ensure the above IDWT and DWT relationship, the following orthogonality condition on the filters $H(\omega)$ and $G(\omega)$ is needed:

$$|H(\omega)|^2 + |G(\omega)|^2 = 1.$$

An example of such $H(\omega)$ and $G(\omega)$ is given by

$$H(\omega) = \frac{1}{2} + \frac{1}{2} e^{-j\omega}$$

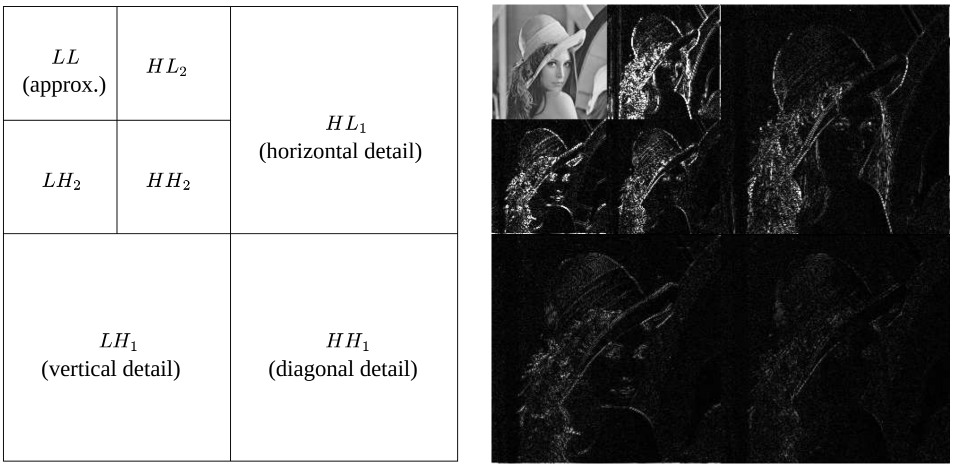

(a) decomposition structure

(b) decomposed image

Figure 2.8: The pyramidal two-level decomposition of an image.

and

$$G(\omega) = \frac{1}{2} - \frac{1}{2}e^{-j\omega},$$

which is known as the Haar wavelet filter. Other common filters used in image processing are the family of Daubechies orthogonal (D-4, D-6, D-8, D-10, D-12) and bi-orthogonal (B-5/3, B-7/9) filters.

The DWT and IDWT for a two dimensional image $F(m, n)$ can be similarly defined by implementing the one dimensional DWT and IDWT for each dimension m and n separately, resulting in the pyramidal representation of an image shown in figure 2.8.

An extension to the dyadic pyramidal decomposition is the wavelet packet decomposition, where the low and high frequency parts are further decomposed every decomposition step. See section **??** for an application of this technique to image watermarking.

2.5.3 Properties of the Wavelet Transform

In figure 2.9, we show two images: Lena on the left and Baboon on the right side. While the Lena image is mostly smooth, except for the feather on the hat, the Baboon image has many textured regions. After a three-level DWT decomposition, we have obtained 10 subbands for each image. The low frequencies (obtained by successive low-pass filtering) are concentrated in the upper-left corner and look like a scaled-down version of the original signal, therefore this subband is also called approximation subband. The high frequency components of the image are represented in the remaining 9 detail subbands.

The wavelet transform has a number of advantages over other transforms, namely the DCT [51]:

(a) Lena

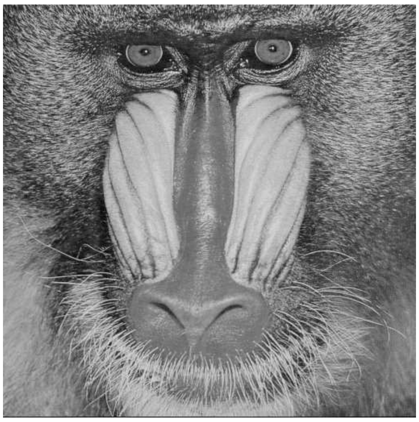

(b) Baboon

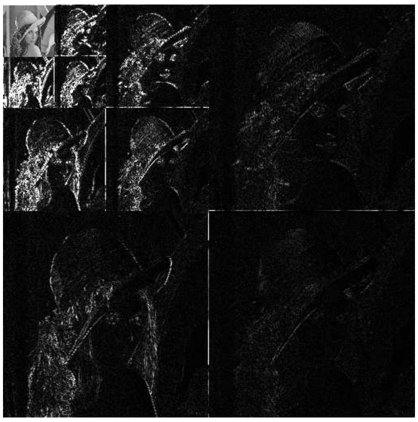

(c) 3-level decomposition

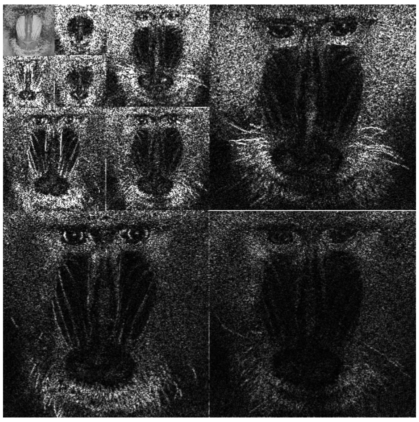

(d) 3-level decomposition

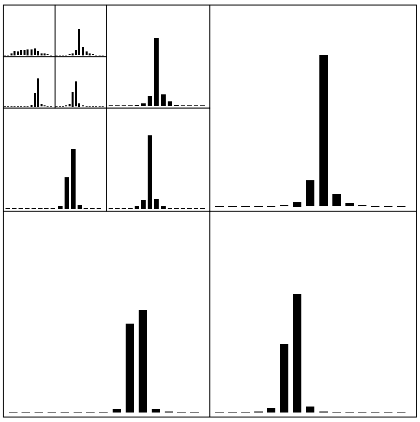

(e) coefficient distribution

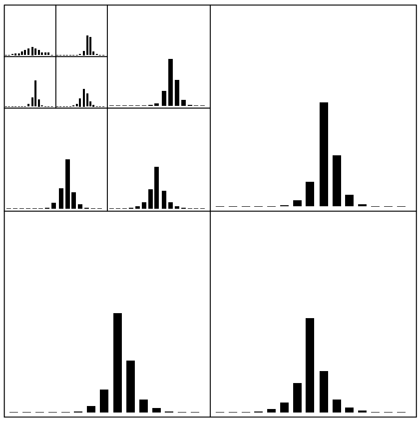

(f) coefficient distribution

Figure 2.9: Image "Lena" (a) is a relatively smooth while image "Baboon" (b) is very textured. The coefficient after a 3-level DWT decomposition are depicted in figure (c) and (d). Note that depending on the orientation of the subband, horizontal, vertical or diagonal image features are emphasized. The distribution of DWT domain coefficients is shown in figure (e) and (f). The smooth image has a more significant peak at the coefficient value 0. The variance is higher in the textured image.

- The wavelet transform is a multi-resolution description of an image: the decoding can be processed sequentially from a low resolution to the higher resolutions.
- The wavelet transform is closer to the human visual system than the DCT. Hence, the artefacts introduced by wavelet domain coding with high compression ratio are less annoying than those introduced at the same bit rate by the DCT.
  Additionally – in the JPEG case – block-shaped artefacts are clearly visible, since image coding based on the DCT usually operates on independent 8×8 blocks.
- The wavelet transform generates a data structure known as scale-space representation. In this image representation, the high frequency signals are precisely located in the pixel domain, while low-frequency signals are precisely located in the frequency domain.
  The spatial resolution of the wavelet transform increases with frequency. Therefore sharp edges which are localized spatially and have a significant high-frequency content, can be seen in the detail subbands and form the contours of the image's objects.While the frequency resolution is independent of the frequency in the DCT domain, it is inversely proportional to frequency in the wavelet domain.

Barni [4], Dugad [55], Xia [266] and other authors identified several advantages which can be exploited by watermarking schemes operating in the wavelet transform domain:

- The hierarchical image representation due to the multi-resolution characteristics of the transform is especially suitable for applications where the image is transmitted progressively, where large amounts of data have to be processed, such as in video application [275], or for real-time systems. Watermarking algorithms that embed a hierarchical or nested watermark can save a lot of computational effort when the mark can be detected early in a progressive transmission. They have to resort to the higher resolution subband only when the watermark could not be detected or extracted from the subbands analyzed previously.
- The wavelet domain allows superior modeling of the human visual system (HVS). It is closer to the hypothetical Cortex transform [257, 153] than the DCT, since it splits the signal into individual bands that can be processed independently. Moreover, the visibility of wavelet quantization noise [258, 259] and the possibilities of visual masking [45] in the wavelet domain have been extensively studied. The HVS is covered in more detail in the next section, 2.6.
- As explained above, the high-resolution subbands allow to locate image features such as edges or textured area easily in the transform domain. Watermarking schemes often put more watermark energy into large DWT coefficients, thus affecting mostly regions, like edges and texture, the HVS is not sensitive to. This is just one example of implicit masking that can be easily exploited in the wavelet domain. Understanding the HVS is

indispensable to achieve imperceptible watermarking with high capacity. For example, Su's [223] watermarking scheme benefits from the locality of the transform coefficients to implement region-of-interest (ROI) coding.

- The wavelet transform is computational efficient and can be implemented in a variety of ways, e.g. by means of filter convolution or via lifting steps [230].

The upcoming next-generation image coding standard, JPEG2000, will be based on the DWT. For some preliminary comments on JPEG2000 and its impact on watermarking see appendix A.

2.6 Human Visual System

The retina of our eye splits a visual signal into different components and each component excites the visual cortex via separate channels [257, 153, 142]. Each component has the following characteristics:

- the spatial location in the image,
- the frequency of the image and
- the orientation of the signal (horizontal, vertical, diagonal).

Based on the knowledge of the structure of the human eye and human visual system (HVS), a hypothetical Cortex transform [257, 153] has been devised that models the known properties. When two signals have similar component characteristics, they excite the same channel in the cortex but are subject to the masking effect. Masking occurs when the detection threshold is increased because of the presence of another stronger signal with similar characteristics. The following effects of the HVS are described in more detail in the next section.

Contrast masking. The detectability of one signal in the presence of another signal.

Frequency sensitivity. The human eye's sensitivity to sine wave gratings at various frequencies.

Luminance sensitivity. The detectability threshold of noise on a constant background.

Just-noticeable-difference (JND) threshold: The threshold beyond which any changes to the respective coefficient will most likely be invisible [262].

A visual model in the frequency domain can therefore be implemented as follows [72, 142]:

1. Apply directional bandpass filters to the host image to obtain the amount of energy the image possesses in each spatial-frequency component.

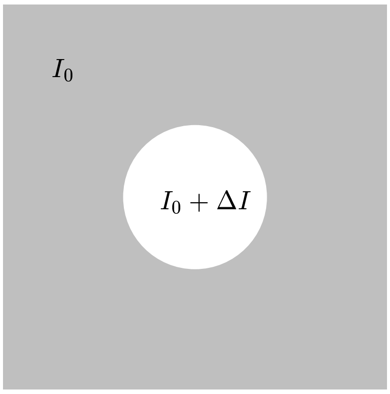

Figure 2.10: Test setup to determine the contrast sensitivity and just noticeable difference ΔI of the human visual system for varying background luminance intensities, I.

   (a) Compute the masking threshold based on the local energy.

2. Scale the watermark energy in each component (assuming the watermark can be decomposed in the same way) so that it is just below this masking threshold.

In order to design optimal digital watermarking methods it is important to take the human visual system (HVS) into account. To understand effects such as masking and contrast sensitivity, a proper model of the visual information processing and representation in the brain is required. By exploiting these phenomena, the performance of watermarking schemes can be greatly improved.

2.6.1 Contrast sensitivity

Contrast sensitivity (also called intensity sensitivity) describes or predicts the visibility of noise. Assuming that the eye is adapted to the luminance of the uniform background I_0, the goal is to determine the minimal difference in luminance ΔI between the central spot and the surrounding area for the human eye to resolve two stimuli. See figure 2.10. This minimal difference is often called just noticeable difference (JND).

Figure 2.11 shows the relationship between the surrounding intensity I and the corresponding minimal contrast, defined as $\Delta I/I$. For intensities in the mid-range, the contrast is approximately constant, while for high and low background intensities the contrast increases, which means that the JND is larger. The approximately constant fraction in the center is called Weber-Fechner fraction, following the Weber law. The Weber-Fechner fraction has been found to be between 1 and 3 %, which means that the JND luminance for the central stimuli is about 0.01 to 0.03 times the surrounding luminance.

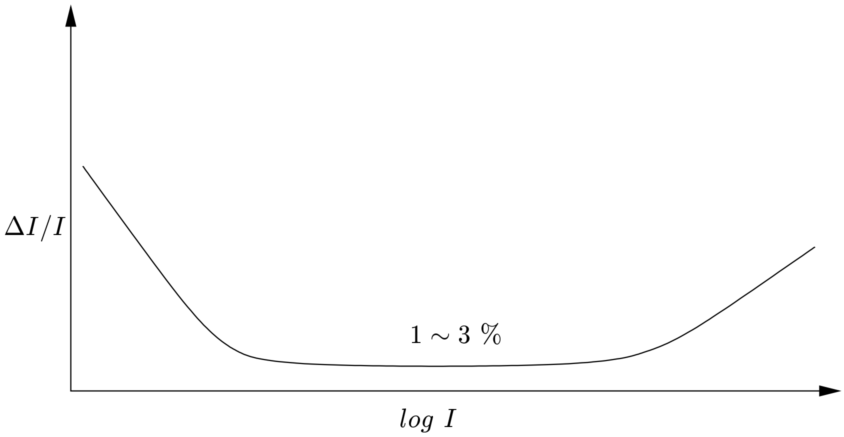

Figure 2.11: The contrast sensitivity function resulting from the test setup presented in figure 2.10. For a large part of the luminance range, the contrast is constant at about 1 to 3 % of the luminance. This fraction is called the Weber-Fechner fraction. For low and high luminance values the sensitivity decreases rapidly.

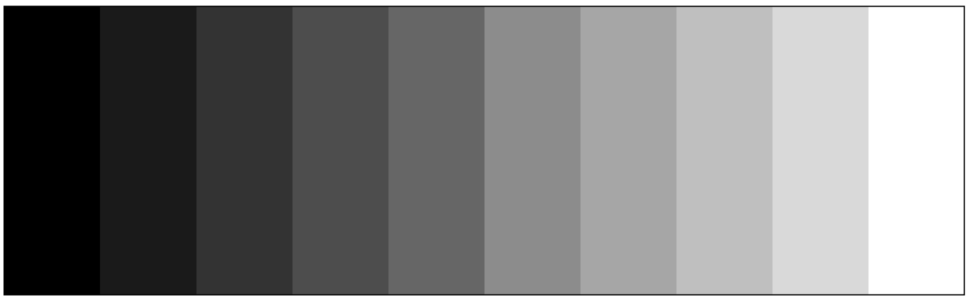

Figure 2.12: The Mach band effect. Although all steps have uniform luminance, we perceive the luminance darker on the right and lighter on the left side of a step.

2.6.2 Spatial frequency sensitivity

Also the spatial frequency (the "shape") has significant influence on the sensitivity of the HVS. The human eye is more sensitive to low-frequency noise. In contrast, high-frequency noise is less visible. The frequency response of the HVS is non-uniform which results in various phenomena, e.g. over-sensitivity in high-contrast areas and especially at edges. One phenomenon called the Mach band effect is illustrated in figure 2.12. The image shows a horizontally varying luminance, where the luminance changes in equal steps. Although each step has a uniform luminance, we perceive the intensity inside a step not as uniform, i.e. brighter on the left and darker on the right side of the edge between two steps.

2.6.3 Masking

Visual masking is a perceptual phenomenon where artefacts are locally masked (hidden) by the image. The image acts as a background signal that reduces the visibility of the artefacts due to image manipulation. Given an image distorted with additive noise, we can observe that the noise is much more visible in flat areas than in areas with high activity, such as textured areas.

2.6.4 Conclusions

According to Kutter [122], the following conclusion can be drawn from our understanding of the HVS with regards to watermarking.

- High frequencies are less visible than low frequencies.
- Studies of the visual cortex showed a multi-resolution characteristic of our visual system. [153]
- In order to embed the watermark as strong as possible, we have to embed it just below the JND. This means, we have to visually adapt the watermark using contrast sensitivity and masking effects.
- The distribution of the blue cones is less dense than the distribution of the red and green cones in the human eye. Therefore, we suggest to put most watermark energy in the blue color component. [125]

2.7 Relationship to Image Compression

Image compression seeks to reduce to number of bits required to represent the image information. Two fundamental concepts of image compression are redundancy reduction and irrelevancy reduction. Redundancy reduction aims at removing duplicate information. Irrelevancy reduction omits part of the information that will not be noticed by the image viewer, namely the HVS (human visual system). There are three types of redundancy:

1. Spatial redundancy or correlation between neighboring pixels.
2. Spectral redundancy or correlation between different frequency bands.
3. Temporal redundancy or correlation between adjacent frames in a sequence of images (in video applications).

Compression technology can be divided into two main groups, lossless and lossy methods. In lossless compression schemes, the reconstructed image, after compression, is numerically identical to the original image. While lossless compression schemes, e.g. JBIG[3], GIF[4] (which is based on LZ-77 coding) or PCX

[3]JBIG ... Joint Bi-level Image expert Group
[4]GIF ... Graphics Interchange Format

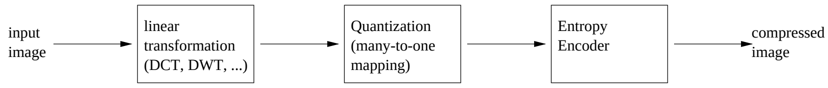

Figure 2.13: Model of a lossy image encoder.

(employing run-length encoding, RLE), can only exploit the redundancy within the image to efficiently code the amount of data, lossy compression methods may discard perceptual insignificant information and, in addition, fall back to efficient lossless coding techniques to represent the remaining salient data. In most schemes, information is discarded in the quantization step, prior to encoding. Due to the discarded information, the compressed image can not be perfectly reconstructed and distortion is introduced into the reconstructed image. The general model of lossy image compression is depicted in figure 2.13.

2.7.1 Distortion Measures

A common measure for compression performance is the achieved compression ratio

$$CR = \frac{\#\, bits\, of\, the\, original\, image}{\#\, bits\, of\, the\, compressed\, image}$$

relative to the distortion. The distortion can be measured as

- the mean squared error (MSE),

  $$MSE = \frac{1}{N}\sum^{N}(F'_i - F_i)^2,$$

  which is the averaged term-by-term difference between the input signal (the original image, F) and the output signal (the watermarked image, F'),

- the signal-to-noise ratio,

  $$SNR = \frac{\frac{1}{N}\sum_i^N F_i^2}{MSE},$$

  which represents the size of the error relative to the input signal – alternatively on a logarithmic scale,

  $$SNR(dB) = 10 \log_{10} SNR,$$

  in units of decibels – or

- the peak-signal-to-noise ratio (PSNR), given by

  $$PSNR(dB) = 10 \log_{10} \frac{F_{peak}^2}{MSE},$$

  where F_{peak} is the peak value of the input signal (usually 255 for 8-bit gray-scale images).

Note, that it is not sufficient to take only quantitative measures into account. More important than any quality metric is the perceptual impression of the human observer.

Lossy compression methods are typically used for natural (photographic) still image compression and can achieve a compression ratios of up to 1 : 100 with acceptable image quality. On the other hand, lossless schemes provide a compression ratio of about 1 : 4 for such images but are better suited for artificial (computer-generated) images (e.g. cartoons, line drawings).

In order to maximize coding efficiency and visual quality, a lossy compression system will exploit the properties of the human visual system (HVS). The human eye is not equally sensitive to image distortion resulting from lossy compression. Therefore, a sophisticated compression system will allow more reconstruction error (e.g. by applying a coarser quantization) in areas than do not have salient image features.

2.7.2 Duality

While lossless compression does not harm a watermarking system in any way (the original data can be perfectly reconstructed), lossy compression methods introduce distortion that has to be taken into account in watermarking applications. Lossy compression techniques are nowadays ubiquitous due to the immediate availability of fast desktop computers on one hand, but limited bandwidth and storage facilities on the other hand. It is therefore imperative to study the effects of lossy image compression on watermarking systems.

More important, the design goal of lossy compression systems is completely contrary to that of watermark embedding systems. The HVS model of the compression system tries to identify and discard perceptually insignificant information of the image. The goal of the watermark system is to embed the watermark information without altering the visual perception of the image. An optimal compression or denoising system would immediately discard any invisible watermark information. This duality has been discussed in the watermarking community. Fortunately, all current compression methods leave enough room for sophisticated watermarking schemes to embed watermark information.

2.7.3 Compression Systems

In the following, we will briefly discuss some common image compression techniques. During the discussion of the watermarking algorithms in section 3 we will re-discover these techniques and see that the compression methodology often forms the basis of the watermark embedding process.

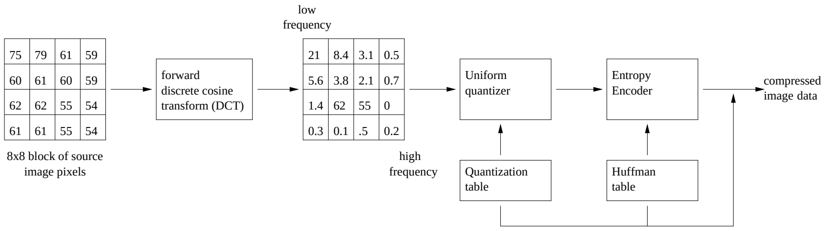

Figure 2.14: JPEG encoder block diagram.

2.7.3.1 JPEG

For still image compression, the JPEG[5] standard has been established by ISO[6] and IEC[7] in 1992 [176]. It is based on the discrete cosine transform (DCT) which can be regarded as a discrete-time version of the Fourier-Cosine series [200]. The baseline JPEG coder essentially operates on a sequence of non-overlapping 8×8 blocks of image samples that are independently transformed to the frequency domain and encoded after quantization. Figure 2.14 shows the key processing steps.

Because adjacent image pixels are highly correlated, most of the signal's energy after the DCT is concentrated in just a few low-frequency coefficients. Next, the transform coefficients are uniformly quantized using a 64-element quantization table which has been designed to preserve visually significant coefficients. In the final step, the quantized coefficients are zig-zag scanned and passed to an entropy coder, usually Huffman coding is employed.

2.7.3.2 Embedded Zerotree Wavelet (EZW) compression

In a wavelet decomposition, as shown in figure 2.15, each coefficient in the high-pass bands of the wavelet transform has four coefficients corresponding to its spatial position in the subband above in frequency. Shapiro named this structure zerotree of wavelet coefficients and presented an elegant algorithm [212] for its efficient coding.

Zerotree coding is based on the hypothesis that if a wavelet coefficient at a coarse scale is insignificant with respect to a given threshold T, then all wavelet coefficients of the same orientation in the same spatial location at a finer scale are likely to be insignificant with respect to T. A zerotree root is encoded with a special symbol indicating that the whole tree is insignificant. This results in gross code symbols savings because at high frequency subbands many insignificant coefficients can be discarded – the tree grows with the power of four per decomposition level.

One of the main advantages of the EZW algorithm is that both, the encoder and the decoder, can terminate the process as soon as the desired target bit rate or distortion rate is met.

[5] JPEG ... Joint Photographic Experts Group
[6] ISO ... International Standards Organization
[7] IEC ... International Electro-Technical Commission

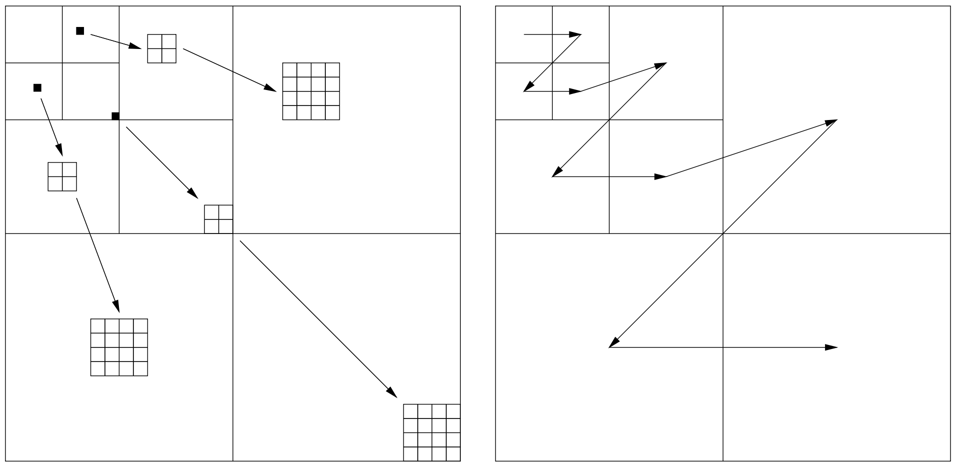

(a) Zerotree structure (b) EZW scan order

Figure 2.15: Zerotree structures (a) and EZW scan-order (b) in a two-level wavelet decomposition.

2.7.3.3 Set Partitioning in Hierarchical Trees (SPIHT) algorithm

The SPIHT algorithm has been proposed by Said [202] as an extension to the EZW compression approach. As opposed to the four different coding symbols of the EZW algorithms, the SPIHT method only transmits a progressive binary representation of the image. SPIHT also builds upon the zerotree structures, but uses a different technique to keep track of the location of significant and insignificant coefficient sets.

Both, EZW and SPIHT, are multi-pass schemes that refine the quality of the transmitted image representation in each coding step. A threshold T is initially set to

$$T_0 = 2^{\lfloor \log_2 c_{max} \rfloor},$$

where c_{max} is the maximum absolute coefficient value in the wavelet transform domain. The initial threshold is subsequently divided by 2,

$$T_i = \frac{T_{i-1}}{2},$$

after each coding step. Therefore, the most significant coefficients are transmitted first, followed by less dominant coefficients as the significance thresholds becomes lower. Variants of the SPIHT algorithm, such as the MTWC [251] approach, calculate the significance threshold per subband.

2.7.3.4 JPEG2000

The current JPEG standard provides excellent compression performance at rates above 0.25 bits per pixel. However, at lower rates, there is a sharp degradation

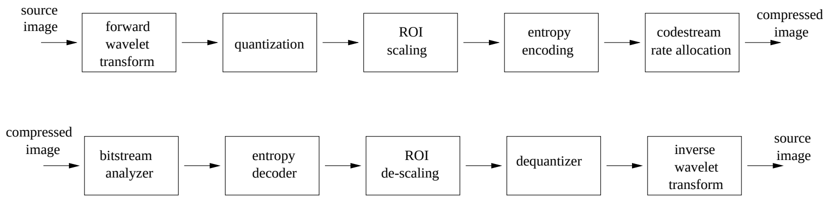

Table 2.2: The JPEG2000 encoder and decoder pipeline.

in the quality of the reconstructed image. This is partly due to the DCT transformation and independent processing of 8×8 blocks, which produces clearly visible block artifacts.

Several other requirements [27], such as progressive transmission, region-of-interest (ROI) coding, scalability, random access, error resilience and lossless as well as lossy operation called for a new standard to overcome the shortcoming of the previous image coding standards: JPEG, P-JPEG (progressive JPEG), JPEG-LS and several optional features on top of "baseline" JPEG. See [205] for a comparative study. Most important with regards to security aspects and watermarking is the inclusion of certain tags in the JPEG2000 header structure that could provide for a security framework. Image security was and is actively discussed during the standardization process. See also appendix A.

The JPEG2000 coding schemes is based on a scheme originally proposed by Taubman [232] and known as EBCOT ("Embedded Block Coding with Optimized Truncation"). The major difference between the previously proposed EZW and SPIHT algorithm is that EBCOT operates on independent, non-overlapping blocks that are coded in several layers to create an embedded, scalable bitstream. Each layer corresponds to a certain distortion level. The partitioning of the available bits between the coding blocks and layers is determined using Lagrangian optimization ("truncation points").

Compared to the JPEG coding scheme, the JPEG2000 standard includes more visual optimization tools, for example based on adaptive frequency weighting and visual masking modeling. [45, 270, 269]

Instead of zerotrees, the EBCOT schemes depends on a per-block quad-tree structure since the independent block coding strategy precludes a tree structure across subbands. These independent blocks are passed down a "coding pipeline" as depicted in figure 2.2.

Chapter 3

Algorithms

Recently, numerous digital watermarking algorithms have been developed to help protect the copyright of digital images and to verify multimedia data integrity. Most watermarking algorithms transform the host image into a domain that facilitates embedding of the watermark information in a robust and imperceptible way. Previous approaches often employed the discrete cosine transform (DCT) to mark perceptually significant coefficients in the low-frequency spectrum [39]. Also, the widely used JPEG compression standard is based on the DCT. However, new requirements such as progressive and low bit-rate transmission, quality scalability and region-of-interest (ROI) coding demand more efficient and versatile image coding. The upcoming compression standard JPEG2000 will be based on the discrete wavelet transform (DWT) to meet the new requirements [27]. Therefore, it is imperative to study watermarking schemes in the wavelet transform domain.

3.1 Classification

In this section, we try to classify the enormous diversity of watermarking approaches, following the classification presented by Loo [142]. Watermarking algorithms can be distinguished in terms of

- the embedding/extraction domain,
  - spatial domain
  - discrete cosine transform (DCT) domain
  - discrete Fourier transform (DFT) domain
  - discrete wavelet transform (DWT) domain
  - miscellaneous domains: e.g. fractal domain, Fourier-Mellin domain, Histogram specification, complex wavelet transform (CWT) domain, ...
- the availability of reference data (e.g. the original host image) for watermark extraction,

  - oblivious (blind)
  - semi-blind
  - non-oblivious (non-blind)

- the host data modification method,
  - linear addition of a spread spectrum signal
  - image fusion (embedding of a "logo")
  - non-linear quantization-and-replace strategy

- the perceptual modeling strategy,
  - no modeling
  - implicit modeling via transform domain properties
  - explicit HVS modeling

- the purpose of the watermarking application,
  - copyright protection, circulation tracking
  - image data verification, image authentication and tamper detection
  - data hiding and image labeling

- and the host media type.
  - still image
  - video
  - special multimedia format such as cartoon, map image

All the algorithms described here are for watermarking monochrome (gray-scale) images. Color images can be dealt with by first transforming them into the YUV color space and then watermarking the luminance component Y. The chrominance components are normally not used because they have a much lower bandwidth or capacity for watermarking purposes. Alternatively, one can watermark each RGB component separately or just watermark a single color component (e.g. the blue component B because of the log HVS sensitivity, see Kutter [125] and Chu [31])

3.2 Overview

This section gives an overview of the proposed watermarking algorithms that operate in the wavelet domain. To the best of our knowledge, it is the most comprehensive compilation of wavelet-based watermarking schemes (December 2000). However, due to the amount of published work in this field (far more than 100 papers per year), we can not guarantee for its completeness.

We organize the algorithms according to their main embedding strategy, additive or quantization, and present them in order of their principal author's name in

Algorithm	Extraction reference data	Modification method	Application	Host media
Barni [4]	blind	additive, HVS	copyright prot.	image
Chae [26]	non-blind	fusion	copyright prot.	image
Chae [25]	blind	quantization	data hiding	image
Chu [31]	blind	quantization	copyright prot.	image
Chen [30]	non-blind	additive	copyright prot.	image
Corvi [34]	non-blind	additive	copyright prot.	image
Davoine [49]	semi-blind	quantization	copyright prot.	image
Dugad [55]	blind	additive	copyright prot.	image
Ejima [65]	blind	quantization	copyright prot.	image
Hsu [81]	non-blind	additive	copyright prot.	image
Inoue [87, 88]	blind/semi-blind	quantization	authentication	image
Jayawardena [90]	blind	quantization	copyright prot.	image
Kanai [99]	blind	additive	copyright prot.	polygons
Kaewkamnerd [92]	blind	additive, HVS	copyright prot.	image
Kim [104]	non-blind	additive	copyright prot.	image
Kim [105]	non-blind	additive, HVS	copyright prot.	image
Kundur [114]	non-blind	fusion, HVS	copyright prot.	image
Kundur [115]	blind	quantization	tamper det.	image
Kundur [118]	blind	quantization	copyright prot.	image
Liang [134]	non-blind	additive	copyright prot.	image
Lin [136]	blind	quantization	tamper detection	image
Loo [143]	non-blind	additive, HVS	copyright prot.	image
Lu [149, 150, 146]	non-blind	additive, HVS	copyright prot.	image
Lu [147]	semi-blind	additive, HVS	copyright prot.	image
Matsui [156]	blind	quantization	data hiding	image
Nicchiotti [166]	non-blind	quantization	copy protection	image
Ohnishi [171]	blind	quantization	copyright prot.	image
O'Ruanaidh [173]	blind	quantization	copyright prot.	image
Pereira [177]	blind	quantization	copyright prot.	image
Podilchuk [189, 190]	non-blind	additive, HVS	copyright prot.	image
Su [223]	non-blind	additive	copyright prot.	image
Tsekeridou [239]	blind	additive	copyright prot.	image
Tzovaras [240]	non-blind	quantization	copyright prot.	image
Vehel [131]	non-blind	quantizaton	copyright prot.	image
Wang [256]	blind/non-blind	additive	copyright prot.	image
Wolfgang [262]	non-blind	additive, HVS	copyright prot.	image
Xia [265, 266]	non-blind	additive	copyright prot.	image
Xie [268, 267]	blind	quantization	authentication	image
Zhu [275]	non-blind	additive	copyright prot.	image/video

Table 3.1: Classification of proposed watermarking algorithms in the wavelet domain (in alphabetical order).

section 3.3 and section 3.4, respectively. Schemes that do not fit any of the above categories are described in section 3.4.4.

In section 3.5, we will summarize the presented algorithms and discuss further concepts that might become important in the near future.

The following principal embedding strategies can be used to embed a watermark in a host image.

1. linear additive embedding,
   (a) Gaussian sequence,
   (b) image fusion
2. non-linear quantization embedding, via
   (a) scalar quantization or
   (b) vector quantization
3. or miscellaneous embedding techniques.

Additive embedding strategies are characterized by the linear modification of the host image and the correlative processing in the detection stage. The quantization schemes on the other hand perform non-linear modifications and detect the embedded message by quantizing the received samples to map them to the nearest reconstruction point.

Before going into the details of the wavelet-based approaches, we want to introduce the concept of additive and quantization watermark embedding by looking at the familiar, now classical, schemes proposed by Cox [39] and Koch [110], both operating in the discrete cosine transform (DCT) domain. For a more complete overview and comparative results of many DCT- and spatial-domain based watermarking schemes, including some early wavelet-based methods, please refer to the work of Jellinek [91].

3.3 Additive Algorithms

3.3.1 Introduction

In additive[1] watermarking algorithms, the signature data is a sequence of numbers w_i of length N that is embedded in a suitably selected subset of the host signal data coefficients, f. The basic and commonly used embedding formula is

$$f'(m,n) = f(m,n)(1 + \alpha \cdot w_i),$$

where α is a weighting factor and f' is the resulting modified host data coefficient carrying the watermark information. Alternative embedding formulas have been proposed by Cox [39], such as

$$f'(m,n) = f(m,n) + \alpha \cdot w_i$$

[1] sometimes also called multiplicative

or, using the logarithm of the original coefficients,

$$f'(m,n) = f(m,n) \cdot e^{\alpha w_i}.$$

An important property of the above formula is that an inverse embedding function,

$$w_i^\star = \frac{f^\star(m,n) - f(m,n)}{\alpha \cdot f(m,n)},$$

can be easily derived to compute $w^\star$ from $f^\star$ given the original host coefficients as a reference. By $f^\star$ we denote the received, possibly altered, image that might contain the watermark w. In the next step, the extracted watermark sequence $w^\star$ is compared to the originally embedded watermark w using the normalized correlation of the sequences as a similarity measure

$$\delta = \frac{w^\star \cdot w}{\|w^\star\| \cdot \|w\|}.$$

The similarity δ varies in the interval $[-1, 1]$; a value well above 0 and close to 1 indicates the extracted sequence $w^\star$ matching the embedded sequence w and therefore we can conclude that the image has been watermarked with w. A detection threshold τ can be established to make the detection decision, $\delta > \tau$. The detection threshold can be derived τ either experimentally by observing the correlation of random sequences (see figure 3.1) or analytically.

For example, a threshold

$$\tau = \frac{\alpha}{S \cdot N} \sum^{N} |f'|$$

can be used, where S, the standard deviation, is 2 or 3.

Of course, the choice of the threshold influences the false-positive and false-negative probability. Hence, a lot of effort has been used to devise reliable methods to compute predictable correlation thresholds and efficient watermark detection systems [5, 160, 137, 78, 7, 172].

The weighting factor α does not necessarily have to be constant over the entire watermark sequence, but can be chosen adaptively to capture local properties of the host signal. This allows to have more energy in the watermark signal and thus have a more robust watermark. For example, certain properties of the human visual system such as masking effects (see section 2.6) can be modeled and exploited.

Before watermark embedding, the host image F is usually subjected to a two-dimensional transform T such as the DCT, DFT or DWT (among others) to derive a frequency representation f of the data, $f = T \times F$. Following the watermarking modifications in the frequency domain, the spatial image representation is regained by applying the inverse transform T^{-1}, $F = T^{-1} \times f$.

Generally, watermarks embedded in the frequency domain have been demonstrated to be more robust to many forms of attacks compared to spatial domain watermarks. In order to achieve robustness, the watermark has to be embedded in salient portions of the host signal. The frequency representation of the host

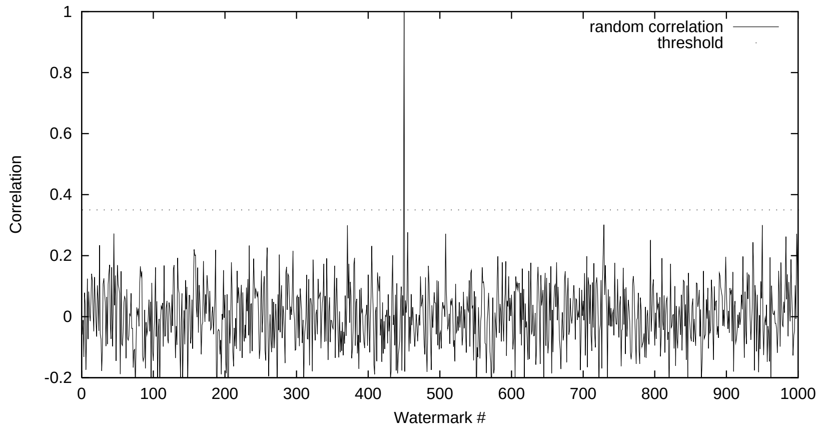

Figure 3.1: Determining the detection threshold experimentally: of the 1000 random sequences tested, only the sequence that was originally embedded yields a high correlation output.

image easily allows to select the low- and mid-frequency coefficients which carry most of the image signal's energy. The selection of suitable transform domain coefficients is one of the most important design issues, as this choice greatly affects robustness, imperceptibility and security of the resulting watermarking scheme. The major difference between the watermarking schemes discussed in the following sections lies in the different coefficient selection strategies.

3.3.1.1 Spread spectrum watermarking

Many watermarking techniques incorporate the ideas of spread spectrum communications to additively embed and extract a pseudo-random noise pattern, either signal-adaptive or non-adaptive. The information bits are spread by simple repetition, error-corrective coding, or some other transform and then modulated with a cryptographically secure pseudo-random noise sequence. A sequence of pseudo-random Gaussian variables is a good model of the noise that is present natural image data. On the other hand, synthetic (computer-generated) images do not contain any noise.

In spread spectrum communications, a narrow-band signal is spread across a wide band of frequencies. This can be accomplished by modulating the narrow-band signal (the watermark information in our case) with a wide-band signal, such as Gaussian noise. The spread watermark signal is similar to the noise already present in the image signal and therefore hard to detect.

3.3.1.2 Generating Gaussian random numbers

Given a source of uniform pseudo-random numbers, the Box-Muller transform can be used to transform uniformly distributed random variables to a new set of random variables with a Gaussian or normal distribution[2]. Algorithm 1

[2] see the notes available at `http://www.taygeta.com/random/gaussian.html`

generates a Gaussian sequence with zero mean and a standard deviation of one, however, it is slow due to the use of trigonometric functions and can have stability problems when `x1` is close to zero.

Algorithm 1 The Box-Muller transform for generating Gaussian distributed random variables from uniformly distributed random variables.

```
double x1 = randf(), x2 = randf();
double y1 = sqrt(-2.0 * ln(x1)) * cos(2.0 * M_PI * x2);
double y2 = sqrt(-2.0 * ln(x1)) * sin(2.0 * M_PI * x2);
```

Algorithm 2, known as the polar form of the Box-Muller transform, overcomes the above weaknesses and is presented in [109]. `randf()` denotes a good, uniform $[0, 1)$ pseudo-random generator.

Algorithm 2 The polar form of the Box-Muller algorithm.

```
double x1, x2, w;
do {
  x1 = 2.0 * ranf() - 1.0;
  x2 = 2.0 * ranf() - 1.0;
  w = x1 * x1 + x2 * x2;
} while ( w >= 1.0 );
w = sqrt((-2.0 * log(w)) / w);
double y1 = mean + x1 * w * deviation;
double y2 = mean + x2 * w * deviation;
```

The sequence that is to be embedded in the host image data can also be binary, i.e. $w_i \in \{-1, 1\}$, or stem from another, smaller image (a "logo"). In the later case, the watermark embedding step is called image fusion. This technique is described in section 3.3.3 in more detail.

3.3.1.3 Information Extraction

One characteristic of basic additive watermark embedding schemes (as proposed by Cox [39]) is that the original image has to be present as a reference for watermark extraction. Furthermore, only information detection is possible via correlation. It is impossible (or at least very time-consuming) to recover the embedded information itself because the embedded information (the signature) would be the seed (usually 32 or 48 bits) of the pseudo-random number generator (PRNG). To retrieve the seed of the PRNG, one would have to try and correlate all possible seed values with the extracted watermark sequence.

To overcome the first problem, i.e. non-blind watermark extraction, several authors [185, 4, 69] presented methods that correlate the watermark sequence w directly with all N coefficients of the received image signal $f^\star$ (mutual correlation), calling upon the Central Limit theorem,

$$\delta = \frac{\sum^N f(m,n)^\star \cdot w_i}{N}$$

and then compare the correlation value δ with some detection threshold τ,

$$\tau = \frac{\alpha}{3 \cdot N} \sum^{N} |f(m,n)^\star| \, .$$

The main drawback of this method is, however, that the image signal itself, has to be treated as noise which makes detection very difficult. In order to reliably detect the presence of a watermark, considerably more coefficients have to be correlated which decreased capacity and robustness. Blind watermark recovery is discussed in more detail in section 3.4, where we present quantization watermarking and a communication model for blind information retrieval.

Semi-blind watermarking is a hybrid approach, where some reference data from the original is available during watermark extraction, but not the original image itself. This can be the watermarked (but otherwise unmodified) image or other salient data facilitating watermarking extraction (such as a model of the probability distribution of the image coefficients).

To address the second problem, i.e. information recovery, several methods have been devised. The problem with correlation and threshold detection is, that it can only detect the presence or absence of a particular watermark sequence. If one is interested in the information bits that are encoded in the watermark sequence, all possible sequences have to be tested – which is clearly computationally infeasible given the random number generator's seed size of at least 32 bits on common systems.

A practical solution proposed by numerous authors [79, 69], is used to encode the information bits in a sequence of real numbers which is approximately Gaussian (see algorithm 3).

3.3.1.4 Example: Algorithm Cox

The most prominent spread spectrum watermarking has been proposed by Cox [39] and is presented here, although operating in the DCT domain, to outline the ancestor of all the algorithms described in the next section.

Authors This algorithm has been developed by Ingemar J. Cox, Joe Kilian, Tom Leighton and Talal G. Shamoon at the NEC Research Institute, Princeton, NJ, USA and is published in [36, 38, 37, 39].

Watermark The mark is a Gaussian sequence of pseudo-random real numbers, length 1000.

Coefficient selection The 1000 largest coefficients in the DCT domain are selected (see figure 3.2).

Embedding Using the additive embedding formula described in section 3.3.

Extraction Using the inverse embedding formula described in section 3.3.

Algorithm 3 Algorithm to encode bits in an approximately Gaussian sequence of real numbers that allows to recover the information easily.

```
void encode_symbol(int N, int b, int s, bits r) {
  bits v = alloc_bits(N + b - 1);
  // gen.  N + b - 1 pseudo-random bits
  for (int i = 0; i < N + b - 1; i++) set_bit(v, i, random() % 2);
  // extract N bits starting from s
  for (i = 0; i < N; i++) set_bit(r, i, get_bit(v, i + s));
  free_bits(v);
}

void encode(char *msg, int n, int N, double m[]) {
  bits r; int i; double mean = 0.0;
  for (i = 0; i < N; i++) m[i] = 0.0; // init.  vector m
  r = alloc_bits(N);
  for (i = 0; i < n; i++) { // encode each symbol of msg.
    encode_symbol(N, 256, msg[i], r); // rep.  symbol as bin.  vector
    for (int j = 0; j < N; j++) m[j] += r[j]; // accum.  vector
  }
  free_bits(r);
  for (i = 0; i < N; i++) mean += m[i]; // calc.  mean
  mean /= (double) N;
  for (i = 0; i < N; i++) m[i] -= mean; // offset seq.  by mean
}

double correlate(double m[], bits v, int N, int s) {
  double c = 0.0;
  for (int i = 0; i < N; i++) c += m[i] * get_bit(v, s + i);
  return c;
}

int decode_symbol(double m[], int N, int b) {
  int i; bits v; int smax = -1; double cmax = 0.0;
  v = alloc_bits(N + b - 1);
  // gen.  N + b - 1 pseudo-random bits
  for (i = 0; i < N + b - 1; i++) set_bit(v, i, random() % 2);
  for (i = 0; i < b; i++)
    if (correlate(m, v, N, i) > cmax)
      smax = i, cmax = correlate(m, v, N, i);
  free_bits(v);
  return smax;
}

void decode(double m[], int N, int n, char *msg) {
  int i; double mean = 0.0;
  for (i = 0; i < N; i++) mean += m[i]; // calc.  mean
  mean /= (double) N;
  for (i = 0; i < N; i++) m[i] -= mean; // offset seq.  by mean
  // decode each symbol of msg.
  for (i = 0; i < n; i++) msg[i] = decode_symbol(m, N, 256);
}
```

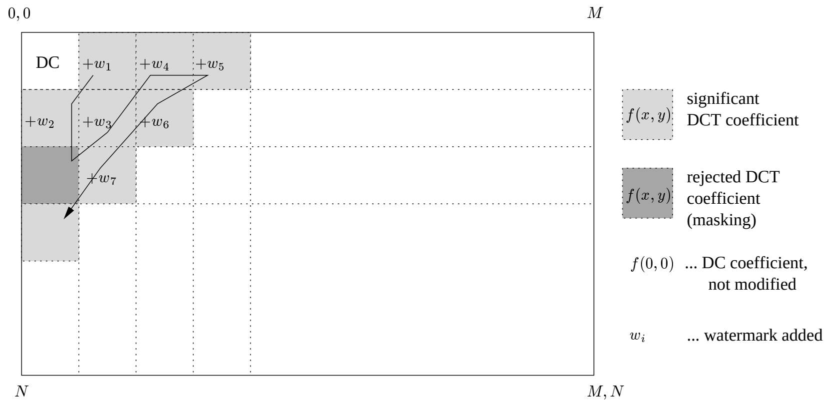

Figure 3.2: Watermarking scheme by Cox.

Discussion The image is first subjected to a global DCT which is computationally expensive. The algorithm achieves good robustness against compression and other signal processing attacks due to the selection of perceptually significant transform domain coefficients. In the papers, the robustness to collusion attacks is discussed, however, the algorithm and many similar schemes is vulnerable to the invertibility attack proposed by Craver [42, 43, 44].

3.3.2 Gaussian-Sequence Algorithms

3.3.2.1 Algorithm Barni

Authors This algorithm has been developed by Mauro Barni, Franco Bartolini, Vito Capellini, Alessandro Lippi and Alessandro Piva at the Department of Information Engineering, University of Siena, Italy and is published in [4].

Watermark The mark is a pseudo-random binary sequence, $w_i \in \{1, -1\}$. The length of that sequence is determined by the dimensions of the host image, M and N respectively, then $i = 0, \ldots 3 \cdot \frac{M}{2} \cdot \frac{N}{2} - 1$.

Decomposition The authors propose using the Daubechies-6 filters to perform a four-level wavelet decomposition. Only detail subband coefficients at the first resolution level are modified to embed the watermark. The coarser detail coefficients are used to calculate a scaling for visual masking.

Coefficient selection All coefficients in the highest resolution detail subbands (LH_1, HL_1, HH_1) are selected.

Embedding Selected coefficients are modified applying the rule for additive embedding described in section 3.3, but taking local image noise sensitivity into account.

$$f'(m,n) = f(m,n) + \alpha \cdot \beta(m,n) \cdot w_i$$

(a) watermarked image

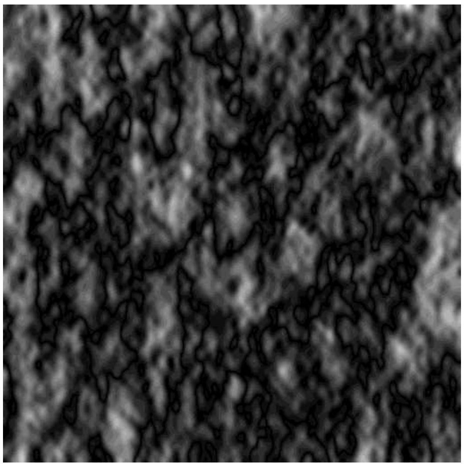

(b) difference image

Parameter	Description
$N = 1000$	length of the pseudo-random sequence
$\alpha = 0.1$	embedding strength

Figure 3.3: Watermarked image (a) and difference image (b), created with Cox's algorithms and the embedding parameters above.

The weighting function $\beta(m, n)$ takes into account the orientation of the subband, i.e. LH, HL, HH, the local brightness based on the corresponding coefficient in the approximation image (LL subband) and the local activity of texture in the neighborhood. The last term combines the local activity in the detail subbands at the coarser level and the local variance of the low-pass subband (LL), both computed in a 2×2 neighborhood corresponding to the location of coefficient $f(m, n)$.

$$\beta(m, n) = \Theta(l, o) \cdot \Lambda(l, m, n) \cdot \Xi(l, m, n),$$

where the first term represents the subband and resolution level sensitivity,

$$\Theta(l, o) = \left\{ \begin{array}{cc} \sqrt{2} & o = HH \\ 1 & otherwise \end{array} \right\} \cdot \left\{ \begin{array}{cc} 1.00 & l = 1 \\ 0.32 & l = 2 \\ 0.16 & l = 3 \\ 0.10 & l = 4 \end{array} \right\},$$

the second term measures local brightness,

$$\Lambda(l, m, n) = \frac{1}{256} f_4^{LL}\left(\frac{m}{2^{4-l}}, \frac{n}{2^{4-l}}\right),$$

and the last term, $\Xi(l, m, n)$, weights the local variance or texture activity.

Extraction The watermark is detected by directly correlating the watermark sequence w_i with the selected image transform coefficients, thus allowing blind detection. See section 3.3.1.3 for details.

Discussion The scheme uses an explicit model of the HVS derived from the problem of coefficient quantization [132]. Each binary watermark value w_i is multiplied, before adding it, by a weighting parameter obtained from the noise sensitivity model. This way, the coefficient is altered to an extent that is just below the visible noticeable difference (JND threshold).

3.3.2.2 Algorithm Corvi

Authors This algorithm has been developed by Marco Corvi and Gianluca Nicchiotti at the Elsag-Bailey Research Department, Genova, Italy and is published in [34, 166].

Watermark The mark is a Gaussian sequence of pseudo-random real numbers, length $32 \times 32 = 1024$.

Decomposition The host image is decomposed to obtain a multi-resolution approximation image of size 32×32.

Coefficient selection All coefficients of the LL subband are selected.

Embedding The watermark is embedded into the approximation image (LL subband) of size 32×32 using the additive embedding formula

$$f'(m,n) = f_{mean} + (f(m,n) - f_{mean}) \cdot (1 + \alpha w_i),$$

where f_{mean} is the average value of the coefficients.

Extraction Using the inverse embedding formula as described in section 3.3.

Discussion The authors state that the DC component of the approximation image is not manipulated since the mean coefficient value f_{mean} is subtracted.
Nicchiotti [166] improves the above scheme with ideas from Nikolaidis's work [167] to achieve non-invertibility. Hence, the coefficients of the approximation image are divided into two subsets using a secret key. The coefficients of one subset are increased by a value K while the coefficients of the other subset are decreased by K. Thus, the mean value of the two subsets is separated by the embedding algorithm. In the detection stage, the algorithm tests if the two coefficient subsets' mean values are apart by approximately $2 \cdot K$.

3.3.2.3 Algorithm Dugad

Authors This algorithm has been developed by Rakesh Dugad, Krishna Ratakonda and Narendra Ahuja at the Department of Electrical and Computer Engineering, University of Illinois, Urbana-Champaign, IL, USA and is published in [55].

Watermark The mark is a Gaussian sequence of pseudo-random real numbers matching size of the detail subbands. Although the watermark is added only to a few selected significant coefficients, using an image size watermark fixes the locations that are manipulated. Hence, there is no dependence on the order of significant coefficients during correlation.

(a) watermarked image

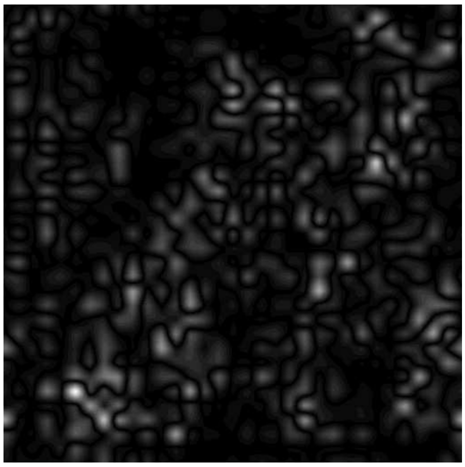

(b) difference image

Parameter	Description
$N = 1000$	length of the pseudo-random sequence
$\alpha = 0.1$	embedding strength

Figure 3.4: Watermarked image (a) and difference image (b), created with Corvi's algorithms and the embedding parameters above.

Decomposition The wavelet transform is a three-level decomposition with Daubechies-8 filters.

Coefficient selection The algorithms selects coefficients in all detail subbands whose magnitude is above a given threshold T_1.

Embedding The equation used for watermark casting in selected significant coefficients is similar to [184], $f'(m,n) = f(m,n) + \alpha \cdot |f(m,n)| \cdot w_i$.

Extraction The blind watermark detection method has already been outlined in section 3.3.1.3. However, only coefficients above the detection threshold $T_2 > T_1$ are considered.

Discussion The author state that visual masking is done implicitly due to the time-frequency localization property of the DWT. Since the detail subbands where the watermark is added contain typically edge information, the signature's energy is concentrated in the edge areas of the image. This makes the watermark invisible because the human eye is less sensitive to texture and edge information.

3.3.2.4 Algorithm J. R. Kim

Authors This algorithm has been developed by Jong Ryul Kim and Young Shik Moon at the Department of Computer Science and Engineering, University of Hanyang, Korea and is published in [104].

(a) watermarked image

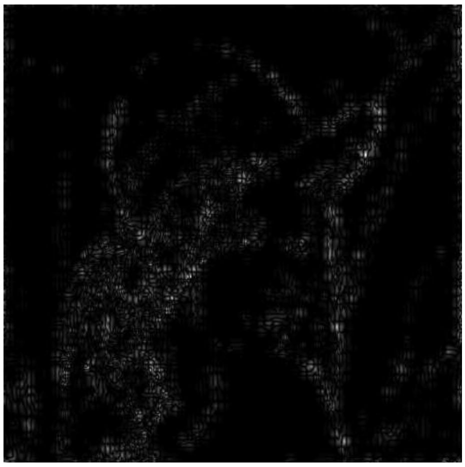

(b) difference image

Parameter	Description
$\alpha = 0.2$	embedding strength
$T_1 = 40.0$	embedding threshold
$T_2 = 50.0$	detection threshold

Figure 3.5: Watermarked image (a) and difference image (b), created with Dugad's algorithm and the embedding parameters above.

Watermark The mark is a Gaussian sequence of pseudo-random real numbers, length 1000.

Decomposition The proposed method uses bi-orthogonal filters to decompose the original image into 3 levels.

Coefficient selection Perceptually significant coefficients are selected using a level-adaptive thresholding scheme. The threshold T_i for decomposition level i depends on the maximum absolute coefficient C_i of all level-i subbands, thus $T_i = 2^{\lfloor log_2 C_i \rfloor - 1}$.

Embedding The additive embedding formula described in section 3.3 is used, however, the scaling factor α is adjusted for each decomposition level. For the LL subband, a scale factor of 0.04 is proposed since the coefficients in the approximation image are generally large. Scale factors of 0.1, 0.2 and 0.4 are used for decomposition level 3, 2 and 1, respectively.

Extraction Using the inverse embedding formula described in section 3.3, but taking the level-adaptive scaling factor into account.

Discussion The proposed algorithms produces a rather robust watermark. The different image modifications due to watermarking in the detail and approximation subbands can be clearly discriminated in the difference image (see figure 3.6). The paper does not address the possibilities of progressive, multi-resolution watermark detection, nor repetitive watermark embedding or watermark weighting to increase robustness.

(a) watermarked image (b) difference image

Parameter	Description
$N = 1000$	length of the pseudo-random sequence
$l = 4$	number of decomposition levels
$\alpha_{LH,HL,HH} = 0.8$	embedding strength for the level 1 detail subband
$\alpha_{LL} = 0.02$	embedding strength for the approximation subband

Figure 3.6: Watermarked image (a) and difference image (b), created with J. R. Kim's algorithm, and the embedding parameters.

3.3.2.5 Algorithm Y. S. Kim

Authors This algorithm has been developed by Young-Sik Kim, O-Hyung Kwon and Rae-Hong Park at the Hanyang University, Korea and is published in [105].

Decomposition The authors propose using a three-level wavelet decomposition.

Watermark The mark is a Gaussian sequence of pseudo-random real numbers, w_i. The length of the sequence in the LL subband is set to 500. In the remaining detail subbands, 4500 coefficients are modified.

Coefficient selection The watermark is added to the large coefficients in each DWT band, except the subbands at the finest resolution level (HL_1, LH_1, HH_1). The number of watermark elements w_i in each of the detail subbands is proportional to the energy of that subband. The energy of a subband, e_s, is defined by

$$e_s = \frac{1}{M \cdot N} \sum_{m=0}^{M-1} \sum_{n=0}^{N-1} f(m,n)^2,$$

where M, N denotes the size of the subband.

Embedding Before embedding, the coefficients are sorted according to their magnitude. Then the watermark is added to the sequence of decreasing wavelet coefficients:

$$f'(m,n) = f(m,n) + \alpha \cdot w_s \cdot f(m,n) \cdot w_i.$$

A relatively small α is used for the LL band, approximately one hundredth of that used for the other subbands. The visual weight w_s is computed per subband and incorporated into the embedding formula to guarantee the invisibility of the watermark.

Extraction Using the inverse embedding formula described in section 3.3.

3.3.2.6 Algorithm Loo

Authors This algorithm has been developed by Patrick Loo and Nick G. Kingsbury at the Department of Engineering, Cambridge University, UK and is published in [143].

Decomposition The host image is decomposed using the dual tree complex wavelet transform (DT-CWT) to obtain a 3-level multi-resolution representation.

Watermark The mark is a bipolar, i.e. $\{-1, 1\}$ pseudo-random bitmap. The bitmap subjected to the CWT before adding to the host image.

Coefficient selection The 1000 largest coefficients in the DCT domain are selected (see figure 3.2).

Embedding The watermark coefficient is scaled and then added to the host image coefficients,

$$f'(m,n) = f(m,n) + \sqrt{\alpha^2 \cdot U(m,n)^2 + \beta^2} \cdot w_i.$$

α and β are level-dependent weights which are designed to embed a strong yet imperceptible watermark. $U(m,n)$ is the average magnitude in a 3×3 neighborhood around the coefficient location.

Extraction Using the inverse embedding formula described in section 3.3.

Discussion The algorithm is based on the dual-tree complex wavelet transform (DT-CWT). This transform has a $2:1$ redundancy for one-dimensional signals and a $4:1$ redundancy in the $2-D$ case. The proposed transform overcomes two drawbacks of the DWT, namely lack of shift invariance and directional selectivity of diagonal features.
Shift invariance means that small shifts in the input signal do not cause major variations in the distribution of energy between wavelet coefficients at different scales. Real DWT filters do not capture the direction of diagonal features. Therefore, the local image activity is not optimally represented, limiting the energy of the signal that can embedded imperceptibly. The DT-CWT overcomes the computational requirements of the undecimated DWT, however, due to the redundancy in the transform domain, some embedded information might be lost the in inverse transform.

3.3.2.7 Algorithm Lu

Authors This algorithm has been developed by Chun-Shien Lu, Hong-Yuan Mark Liao, Shih-Kun Huang and Chwen-Jye Sze at the Institute of Information Science, Academia Sinica, Taipei, Taiwan and is published in [149, 150, 147, 146].

Watermark The mark is a pseudo-random Gaussian sequence of real numbers [150], a bipolar sequence [149], i.e. $w_i \in \{-1, 1\}$, or a gray-scale image [146] matching the number of selected coefficients. Half of the watermark is positively modulated, the other half is negatively modulated.

Decomposition Using a 3-level wavelet transform.

Coefficient selection A wavelet coefficient is selected for modulation if the magnitude is larger than the corresponding JND threshold [259]. Since two complementary watermarks are embedded, the locations for the two watermarks are interleaved.

Embedding Prior to the actual embedding process, the wavelet coefficients are sorted in increasing order based on their magnitude. The Gaussian watermark sequence is sorted as well. Each time, a pair of wavelet coefficients, $(f_{positive}, f_{negative})$, is fetched from the top of the sorted host image coefficient sequence, f, and a pair of watermark values, (w_{top}, w_{bottom}), is fetched the top and the bottom of the sorted watermark sequence, w. The following modulation rules for positive modulation,

$$f' = \begin{cases} f_{positive} + J \cdot w_{bottom} \cdot \alpha, & f_{positive \geq 0} \\ f_{positive} + J \cdot w_{top} \cdot \alpha & f_{positive < 0} \end{cases}$$

and negative modulation,

$$f' = \begin{cases} f_{negative} + J \cdot w_{top} \cdot \alpha & f_{negative \geq 0} \\ f_{positive} + J \cdot w_{bottom} \cdot \alpha & f_{negative < 0} \end{cases}$$

are applied to the selected wavelet coefficients to adaptively embed the watermark. J denotes the JND value of the selected wavelet coefficient based on the visual model [259, 262, 190]. The weight α controls the maximum possible modification and is chosen differently for the high- and low-frequency subbands (i.e. approximation and detail subbands).

Extraction Watermark extraction is achieved re-ordering the transform coefficients and applying the inverse formula,

$$w^\star = \frac{f^\star - f}{J - \alpha}.$$

Discussion The embedding and extraction method in this algorithm is based on [190, 262], described in section 3.3.2.9.
The complementary modulation approach proposed in this work is not limited to the wavelet domain but can be applied to all spread spectrum watermarking schemes. It is assumed that an arbitrary attack decreases ***or*** increases the majority, that is significantly more than 50 percent, of the

transformed coefficients. The chance that an attack will make the number of increased and the number of decreased coefficients equal is very small. By simultaneously embedding two complimentary watermarks, one of the two marks will be significantly stronger after attack, thus performing better than random modulation.
The authors claim that sorting the wavelet coefficients ("relocation strategy") before embedding and extraction improves robustness to asynchronous phenomena such as jitter [179] or StirMark attacks [178].

3.3.2.8 Algorithm Lu (blind)

This algorithm [147] is an variation of the methods described in section 3.3.2.7 which allows semi-blind watermark extraction. The original and the received image is modeled during the extraction step using a generalized Gaussian model of the wavelet coefficients. Hence, the original image is not needed for watermark extraction – just a set of image-dependent parameters that describe the wavelet coefficient probability distribution has to be transmitted. Only the high-frequency bands can be accurately modeled, therefore the host image coefficient selection is limited to certain detail subbands.

3.3.2.9 Algorithm Podilchuk/Wolfgang

Authors This algorithm has been developed by Christine I. Podilchuk at Bell Laboratories, Lucent Technologies, Murray Hill, NJ, USA, Wenjun Zeng at Sharp Laboratories of America, Inc., Camas, WA, USA, Raymond B. Wolfgang and Edward J. Delp at the Video and Image Processing Laboratory, School of Electrical and Computer Engineering, Purdue University, West Lafayette, IN, USA and is published in [189, 190, 262, 261].

Watermark The mark is a Gaussian sequence of pseudo-random real numbers, depending on the image's capacity (according to the visual model).

Decomposition The schemes is proposed for a four-level wavelet decomposition using 7/9-bi-orthogonal filters.

Coefficient selection Only transform coefficients $f(m,n)$ above their corresponding JND threshold $j(m,n)$ are selected. The selection is not limited to perceptual significant parts of the image (such as, say, the first 1000 DCT coefficients as in Cox's [39] scheme).

Embedding Using the additive embedding formula described in section 3.3, taking the JND threshold into account:

$$f'(m,n) = \begin{cases} f(m,n) + j(m,n) \cdot w_i & if\, f(m,n) > j(m,n) \\ f(m,n) & otherwise \end{cases} .$$

The JND visual model employed in the proposed watermarking algorithm is based on the work of Watson [258, 259].

Extraction The watermark is extracted with the help of the original image using the methods proposed in section 3.3. Before correlating the coefficients with the watermark sequence, a filtering step is performed which rejects all coefficients below the current JND threshold.
For the wavelet-based variation of this algorithm, the correlation calculation is performed separately for each multi-resolution level and the peak correlation is considered. For example, cropping the image will impact the watermark values more in the lower frequency levels because the watermark values in the higher level benefit from smaller spatial support of the mark. On the other hand, low-pass filtering operations affect mostly the high-level coefficients.

Discussion The authors build one the work of Cox [39] and add a scaling factor in the embedding formula that depends on the signal strength of a particular frequency component. This weighting factor is derived from a visual model and based on the JND paradigm (see section 2.6). The JND approach can be easily incorporated into the watermarking scheme and provides an upper bound on the watermark intensity. Hence, an imperceptible yet robust watermark can be embedded. It is important to note that the visual modeling is much simpler in the DWT than in DCT domain.
The proposed scheme can be applied to DCT-blocks as well as DWT coefficients. In the DCT domain, the scheme has the advantage of encoding the watermark in the JPEG bitstream while in the DWT domain, the schemes is compatible with the JPEG2000 standard [271].
In [261], the authors compare the robustness of watermarks embedded in the DCT versus the DWT domain – especially when subjected to lossy compression. In their experiments with the JPEG and the wavelet-based EZW [212] coder, they found that it is beneficial to match the compression and watermarking domain, in particular for high compression rates.

3.3.2.10 Algorithm Xia

Authors This algorithm has been developed by Xiang-Gen Xia, Charles G. Boncelet and Gonzalo R. Arce at the Department of Electrical and Computer Engineering, University of Delaware, Newark, DE, USA and is published in [265, 266].

Watermark The mark is a Gaussian sequence of pseudo-random real numbers.

Decomposition The authors propose using a 2-level decomposition and the Haar wavelet filter.

Coefficient selection The watermark is embedded in large coefficients of the high and middle frequency bands (detail subbands). The LL subband does not carry any watermark information.

Embedding Using the additive embedding formula

$$f(m,n)' = f(m,n) + \alpha \cdot f(m,n)^{beta} \cdot w_i,$$

where α is the embedding strength and b indicates the amplification of large coefficients.

(a) watermarked image

(b) difference image

Parameter	Description
$\alpha = 0.2$	embedding strength
$\beta = 1.2$	coefficient amplification

Figure 3.7: Watermarked image (a) and difference image (b), created with Xia's algorithm and the above embedding parameters.

Extraction Using the inverse embedding formula, very similar to that described in section 3.3.

Discussion The authors discuss the advantage of a multi-resolution representation. The detection process can benefit from the hierarchical decomposition and save unnecessary computations if the watermark can already be detected in an early stage of the decomposition.
Since large coefficients in the detail subbands generally indicate edges, this algorithms places most watermark energy in areas containing edges and texture. This implicit masking effect can be seen in the difference image of figure 3.7. The human eye is less sensitive to changes in edge and texture information, as opposed to changes in low-frequency components of the signal that are concentrated in the LL subband of the transform.
The authors claim that the DWT has advantages over the DCT after rescaling attacks. The DCT coefficients of the rescaled image are shifted in two directions from the locations of the original image. Due to the localization of the DWT, not only in the time but also in the frequency domain, the correlation does not suffer as much as in the DCT case.
In the later paper [266], a correlation method is described that uses the peak correlation at all offsets γ,

$$\delta = \max_{\gamma} \frac{\sum_i^N w_i^\star \cdot w_{(i+\gamma)\, mod\, N}}{\|w^\star\| \, \|w\|}$$

3.3.2.11 Algorithm Wang

Authors This algorithm has been developed by Houng-Jyh Mike Wang, Po-Chyi Su and C.-C. Jay Kuo at the Department of Electrical Engineering, University of Southern California, Los Angeles, CA, USA and is published in [256, 253, 252, 254, 255, 223, 222].

Watermark The mark is a Gaussian sequence of pseudo-random real numbers matching the number of selected coefficients.

Decomposition The authors do not specify how to decompose the image, however, it seems plausible to perform at least 5 decomposition steps.

Coefficient selection The watermark is added to significant coefficients in significant subbands. The search for these locations is based on the design principles of the multi-threshold wavelet coder (MTWC) [251, 250], namely successive subband quantization (SSQ) and bit-plane coding.
The algorithm selects coefficients whose magnitude is larger than the current subband threshold, $T_{s,i}$. After watermarking a subband, the subband's threshold is divided by 2. The initial threshold of a subband s is determined by

$$T_{s,0} = \beta_s \frac{\max|f_s|}{2}.$$

b_s is used to weight the subbands.
Algorithm 4 starts with the most significant subband (with the highest initial threshold $T_{s,0}$) and proceeds until enough coefficients are selected. A coefficient is selected only once for embedding, however, a subband may be visited multiple times. Only detail subbands are considered for watermark casting, the approximation subband is not selected.

Algorithm 4 Significant subband and coefficient selection algorithm in Wang's watermarking scheme.

```
for each subband s {
  calculate initial threshold T_{s,0}
  set all subband coefficients unselected
}

while more watermark symbols to cast {
  s_max is the subband with maximum value of T_s
  for each coefficient c_i in s_max {
    if c_i is unselected and |c_i| > T_{s_max} {
      select c_i for watermark casting
     mark c_i selected
    }
  }
  update threshold of s_max : T_{s_max}, = T_{s_max}/2
}
```

Embedding Using the additive embedding formula

$$f_s(m,n)' = f_s(m,n) + \alpha_s \cdot T_s \cdot w_i,$$

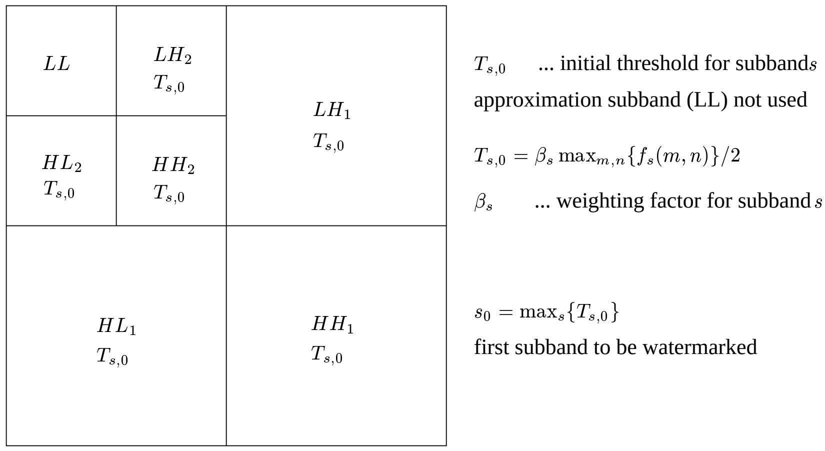

Figure 3.8: Watermarking scheme by Wang.

where α_s is the scaling factors for the subband s.

Extraction Using the inverse embedding formula, very similar to that described in section 3.3.

Discussion In one of the later papers [254], the authors focus more on security aspects and propose to use a key-dependent coefficient skipping scheme to achieve non-invertibility and, optionally, a key-dependent transform structure in order to conceal the embedding locations.
By weighting the wavelet coefficients according to their perceptual importance via the selection thresholds T_s, the resulting image distortion can be constrained to an acceptable fidelity loss. Values for the subband weighting factor b_s can be found in the work of Barni [4] discussed in section 3.3.2.1.
For cartoon and map image data, Su [222] modifies the embedding method to allow bitmap embedding. A selected subband is divided into several blocks of equal size. Each block carries one bit of watermark information (the bitmap).
In [223], Su extends this scheme for image labeling and region of interest (ROI) watermarking applications. A binary ROI map is constructed for each resolution level to select the region to be watermarked. To this end, the spatial ROI mask has to be scaled to fit the multi-resolution subbands of the wavelet transform domain. Repeatedly, the dominant value of each 2×2 ROI block is mapped to the coefficient in the next coarser map.

3.3.2.12 Algorithm Wang (blind)

Based on Wang's scheme discussed in section 3.3.2.11, a variation of the algorithm has been proposed that allows blind detection. Since the original image is not available during the watermark extraction step, the original coefficients have

(a) watermarked image

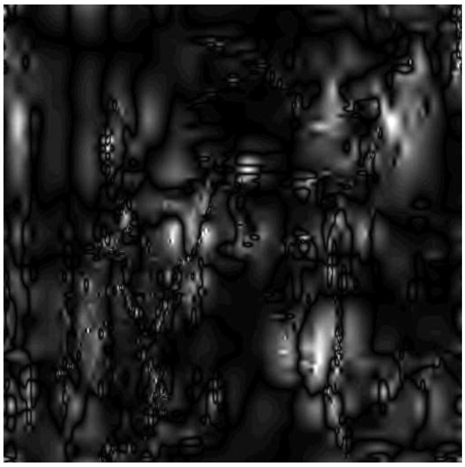

(b) difference image

Parameter	Description
$N = 1000$	watermark sequence length
$\alpha = 0.3$	embedding strength
$\beta = 1.0$	subband weighting factor (not used)

Figure 3.9: Watermarked image (a) and difference (b) image, created with Wang's algorithm and the embedding parameters above.

to be modeled and estimated. Therefore, the blind embedding and extraction algorithms rely on truncating (quantizing) selected coefficients to well-defined values.

Let $f_s(m,n)$ be a selected coefficient in subband s, i.e. $T_s < |f_s(m,n)| \leq 2 \cdot T_s$. Then, the coefficient is modified according to

$$f'_s(m,n) = sign \cdot \Delta_p(|f_s(m,n|) + \alpha_s T_s w_i,$$

where $sign$ is the sign value of the coefficient $f_s(m,n)$ and the quantization operation Δ_p is defined as

$$\Delta_p(x) = (1 + 2 \cdot p \cdot \alpha_s) \cdot T_s.$$

The integer value p is chosen such that the distance

$$|\Delta_p(|f_s(m,n)|) - |f_s(m,n)||$$

between quantized and original coefficient is minimal.

In Wang's blind watermark extraction formula, the term representing the original image coefficient is replaced with the approximation

$$sign \cdot \Delta_p(|f_s^\star(m,n)|),$$

thus we get

$$w_i^\star = sign \cdot \Delta_p(|f_s^\star(m,n)|) - f_s^\star(m,n).$$

Analysis of the blind detection method shows that the detection performance is four times weaker compared with the non-blind scheme for identical values of α (same acceptable distortion).

The authors discuss a security vulnerability of their blind scheme in [254]. Since the coefficient selection algorithm depends on the correct order of significant subbands which is determined by the subband's maximum coefficient magnitude, the watermark detection scheme can be easily foiled by manipulating the subband's maximum coefficient. The authors show that tweaking the subband's maximum coefficient add only little distortion to the image.

3.3.2.13 Algorithm Zhu

Authors This algorithm has been jointly developed by Wenwu Zhu at the Bell Labs, Lucent, Technology, Homdel, NJ, USA and Zixiang Xiong and Ya-Qin Zhang, both with the Department of Electrical Engineering, University of Hawaii, Honolulu, HI, USA and is published in [275, 276].

Watermark The mark is a Gaussian sequence of pseudo-random real numbers. The length of the watermark sequence equals the number of detail coefficients.

Decomposition The authors suggest a four-level wavelet decomposition.

Coefficient selection All high-pass subband coefficients are selected, sparing only the LL subband.

Embedding Using the additive embedding formula described in section 3.3. The watermark sequence at different resolution levels is nested,

$$\ldots \subset W_3 \subset W_2 \subset W_1,$$

where W_j denotes the watermark sequence w_i at resolution level j. The length of W_j is given by

$$N_j = 3 \cdot \frac{M^2}{2^{2 \cdot j}}.$$

Extraction Using the inverse embedding formula described in section 3.3. The peak correlation over all resolution levels is used to detect the presence or absence of the watermark.

Discussion Because of its simple structure, the algorithm can be easily incorporated in video watermarking application based on a 3-D wavelet transform.

3.3.3 Image-Fusion Algorithms

Watermarking algorithms which embed meaningful data in form of a logo image instead of a pseudo-random number sequence are called image-fusion watermarking algorithms. The logo image is generally smaller than the host image. Before being added to the host signal, the logo image is encrypted (decorrelated) [247, 248, 206] and suitably transformed.

(a) watermarked image

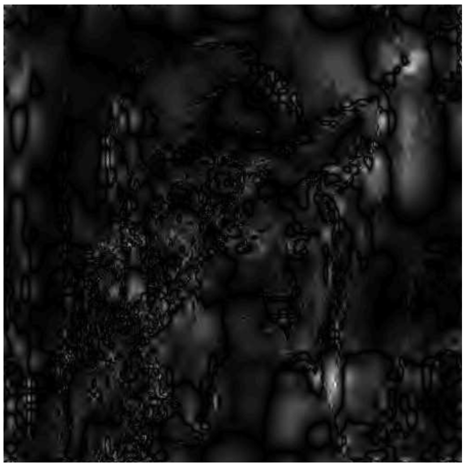

(b) difference image

Parameter	Description
$\alpha = 0.2$	embedding strength
$L = 7$	decomposition level

Figure 3.10: Watermarked image (a) and difference image (b), created with Zhu's algorithm and the embedding parameters above.

There are two important advantages of embedding a logo image as watermark data. First, the extracted image can be correlated with the originally embedded image by a human observer, building on the superior pattern-matching capabilities of the human brain. Second, the existence of a visual logo in the questionable image might be much better proof of ownership than a high statistical correlation value. Fortunately, the transition from Gaussian-sequence watermarking to logo image watermarking is not very difficult.

3.3.3.1 Algorithm Chae

Authors This algorithm has been developed by Jong Jin Chae and B. S. Manjunath at the Department of Electrical and Computer Engineering, University of California, Santa Barbara, CA, USA and is published in [26, 21].

Watermark The watermark is a gray scale image, with as much as 25% of the host image size.

Decomposition The authors propose using a 1-level decomposition on both, the host and and the logo image, with the Haar wavelet filter. The wavelet domain representation of the host image is denoted by $f(m,n)$, the DWT coefficients of the logo image by $w(m,n)$.

Coefficient selection Each coefficient of the host image is modified to embed the logo image.

(a) host image (b) logo image (c) fused image

Figure 3.11: The image fusion process: the host image (a) is fused with the logo image (b), a quarter of the size of the host image. Chae's algorithm produces the fused (combined) image (c), the embedding parameter α set to 0.5 to make the fusion visible.

Embedding The embedding process is depicted in figure 3.12. First, the host and logo image coefficients of each subband are linearly scaled to 24 bits per coefficient (excluding the sign bit). Since the logo image is a quarter of the size of the host image, the coefficients have to be expanded. Let A, B, C represent, respectively, the most significant byte (MSB), the middle byte, and the least significant byte (LSB) of the 24 bit representation of a logo coefficient. Three 24-bit numbers A', B', C' are generated with their most significant byte set to A, B, C, respectively, and with their two least significant bytes set to zero. Then a 2×2 expanded block is formed as shown in the figure.

After adding the expanded logo image to a scaled version of the host image,

$$f'(m,n) = \alpha f(m,n) + w(m,n),$$

the 24-bit representation is scaled back using the original minimal and maximal coefficient values per subband. Finally, the fused (combined) image is produced via the IDWT.

Extraction Using the inverse embedding formula, very similar to that described in section 3.3.

Discussion The proposed method allows to hide surprisingly high amounts of image data in a host image. The current implementation is limited to logo images that are a quarter of the size of the host image. However, this constraint can easily be removed by exploiting the multi-resolution property of the wavelet transform and performing more decomposition steps.

Chae's scheme hides most of the logo image's energy in the low frequency subband of the host image.

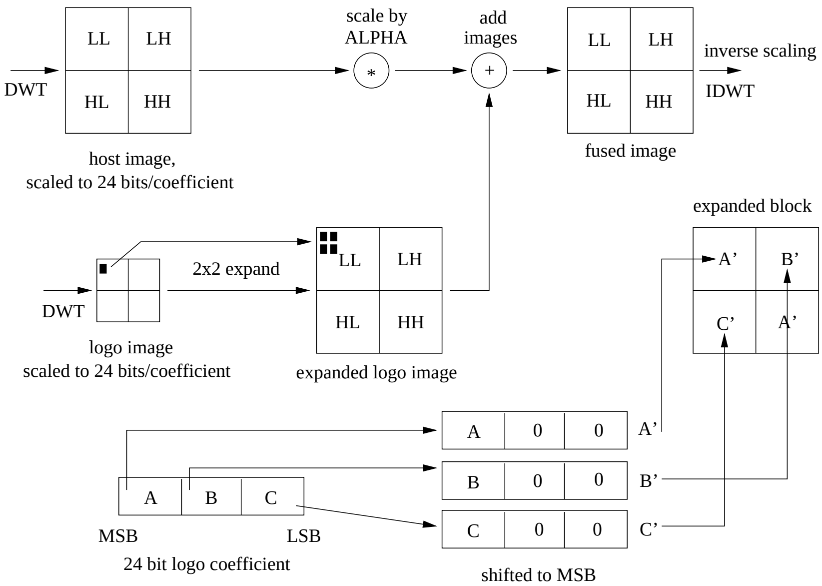

Figure 3.12: Watermarking scheme by Chae.

(a) watermarked image

(b) difference image

Parameter	Description
$\alpha = 15$	embedding factor

Figure 3.13: Watermarked image (a) and difference image (b), created with Chae's algorithm and the embedding parameters above.

3.3.3.2 Algorithm Kundur (fusion)

Authors This algorithm has been developed by Deepa Kundur and Dimitrios Hatzinakos at the Department of Electrical and Computer Engineering, University of Toronto, ON, Canada and is published in [114].

Watermark The mark is a logo image which is decomposed using the DWT. Instead of an image, any noise-like two-dimensional array of binary or real numbers can be used. The watermark is assumed to be a factor of 2^M smaller than the host image and have dimensions $2 \cdot N_w$ and $2 \cdot M_w$, respectively.

Decomposition Both, the host image and the watermark data, is transformed into the wavelet domain. The host image is decomposed in L steps where L is an integer less or equal to M.

Coefficient selection The watermark is embedded in all detail subbands.

Embedding The detail images of the host at each resolution level are segmented into non-overlapping block of size $N_w \times M_w$. The blocks are denoted by $f^i_{k,l}(m,n)$, where $i = 1, \ldots, 2^{2\cdot(M-l)}$. k and l are the orientation and the resolution level, respectively, of the underlying subband. The watermark is embedded by simple scaled addition of the watermark to the particular $N_w \times M_w$ detail component of the host image,

$$f^i_{k,l}(m,n) = f^i_{k,l}(m,n) + \alpha_{k,l} \cdot \sqrt{S(f^i_{k,l}(m,n))} \cdot w_{k,1}(m,n).$$

The salience S of a localized block is computed and as used as one of the scaling factors in the embedding formula,

$$S(f^i_{k,l}(m,n)) = \sum_{u,v} C(u,v) \cdot \left|T(f^i_{k,l}(m,n)\right|^2.$$

$C(u,v)$ is the contrast sensitivity matrix according to Dooley and T is the discrete Fourier transform. The parameter $\alpha_{k,l}$ controls the visibility versus robustness of the embedded watermark.

Extraction Using the inverse embedding formula, very similar to that described in section 3.3.

Discussion The algorithm uses a rather complex explicit model of the HVS. The paper provides rules for choosing all parameters of the HVS model and the scaling parameters.
If a logo image was embedded, either statistical correlation or visual detection can be employed to verify the presence of the watermark.
The authors claim that the wavelet domain representation of an image contains the image components in bands of approximately equal bandwidth on a logarithmic scale, much like the HVS splits an image into several components. It is therefore expected that the DWT will allow the independent processing of the separate components like the human eye. This property makes the wavelet decomposition very popular for image fusion applications.

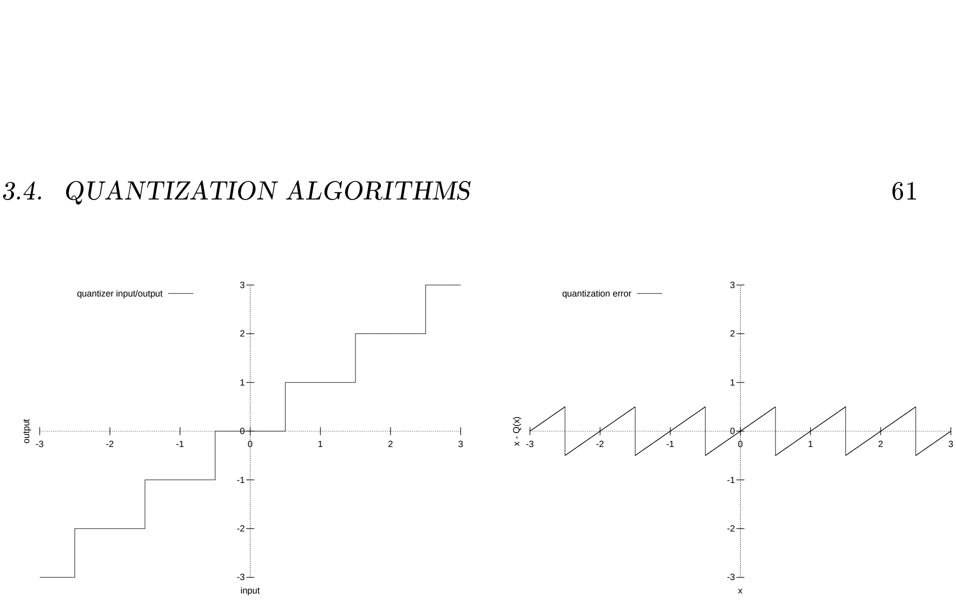

(a) input/output map (b) quantization error

Figure 3.14: Quantizer input/output map (a) and quantization error (b) of the `rint()` function.

3.3.3.3 Algorithm Lu (fusion)

Based on Lu's algorithm described in section 3.3.2.7, there also an image fusion methods has been proposed [146] that modifies the approximation and detail subbands according to the "cocktail modulation" technique.

3.4 Quantization Algorithms

3.4.1 Introduction

The process of mapping a large – possibly infinite – set of values to a much smaller set is called quantization. Since quantization reduces the number of distinct symbols that have to be coded, it is central to many lossy compression schemes (see section 2.7).

We distinguish between scalar and vector quantization. In the first case, the quantizer takes and outputs scalar values, while in the later case, the quantizer operates on vectors. Watermarking schemes based on these quantization techniques are describes in section 3.4.2 and section 3.4.3, respectively.

A quantizer consists of two mappings: an encoder mapping and a decoder mapping. The encoder divides the range of source values into a number of intervals. Each interval is represented by a codeword. The encoder represents all the source values that fall into a particular interval by the codeword assigned to that interval. As there could be many – possibly infinitely many – distinct samples that can fall in any given interval, the encoder mapping is irreversible. For every codeword generated by the encoder, the decoder generates a reconstruction value. An example for a simple quantizer would be the `rint()` function found in the C standard library which maps a real number to the nearest integer value, see figure 3.14 for an illustration.

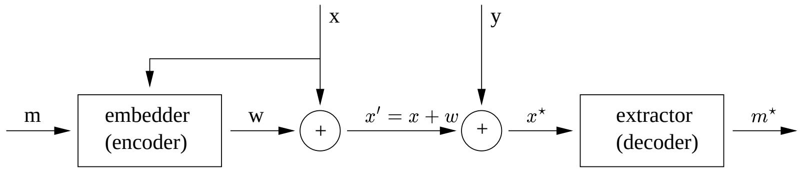

Figure 3.15: Communication model where the original signal is not available at the receiver's side.

3.4.1.1 Host-signal interference and Self-noise suppression

Blind digital watermarking is the communication of information via multimedia host data (e.g. digital images) where the unmodified host data is not available to the watermark detector [63]. Many blind watermarking schemes, especially from those based on the additive embedding strategy (see section 3.3), suffer from host-signal interference when correlating the watermark sequence with received image data.

For additive Gaussian noise attacks, Chen showed using the theoretic results of Costa [35] that interference from the host signal can be eliminated when embedding the watermark in a non-linear way, e.g. via quantization. Thus, the theoretical capacity of a blind watermarking scheme is equal to methods where the receiver has access to the host signal. A model of this communication problem is depicted in figure 3.15.

The message m is to be transmitted with a power constraint (to ensure imperceptibility). The interfering Gaussian noise sources, x (the host image) and the processing noise p (assumed to be Gaussian as well), are not known to the decoder. However, the encoder has knowledge (side information [40]) of x. The decoder must be able to decode the watermark message m from the received composite signal $x^\star$ without having access to the original host signal x.

The communication channel has two sources of noise – x, the noise due to the original image, and p, the noise due to image processing, compression, watermark attacks.

Costa's solution for the blind watermarking problem is not practical since a huge codebook is involved. Therefore, simpler quantization methods have been derived, such as lattice-structured codes and other quantization-based modulation schemes.

3.4.1.2 Image Decomposition and Energy distribution

For most image decompositions, such as the DCT, DWT, DWT, the low frequency bands (or channels) carry a large amount of the image's energy, and thus represent the majority of the image-noise or self-noise [197, 198]. On the other hand, these bands are hardly affected by common image processing operations and therefore contribute only little processing noise. The high-frequency bands, which suffer most from image processing noise, contribute only marginal amounts of self-noise.

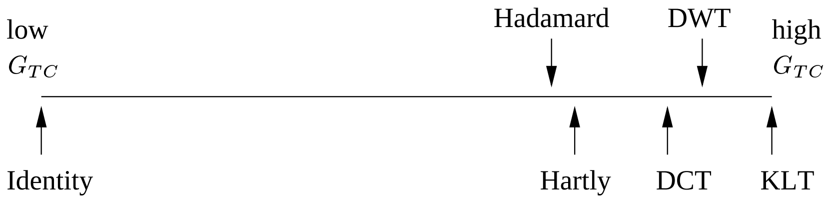

Table 3.2: Different transforms on the energy compaction scale. G_{TC} is the transform coding gain.

For non-blind watermarking applications, image noise is not an issue because it can be canceled in the watermark extraction stage by simply subtracting the original image from the received image. This principle has motivated the choice of e.g. low-frequency DCT coefficients in Cox's scheme [39] and lead to robust watermarking since image processing and most attacks are confined to the high-frequency coefficients that do not carry watermark information.

If the original image is not available in the watermark communication process, i.e. blind watermarking, embedding the watermark in the low frequency components of the image is a problematic choice due to the predominant self-noise. Generally, blind additive watermarking methods favor the mid-frequency regions for this reason. At this point, one might consider using a transform that does not have as much coding gain, G_{TC}, [51] or allow for less energy compaction [198]. Figure 3.2 shows the position of different transforms on the "scale" of energy compaction. At the left end is the Identity transform, on the extreme right we have the Karhunen-Loéve transform (KLT), which shows the best energy compaction ability.

Contrary to linear methods, non-linear watermark embedding strategies are capable of utilizing the low-frequency bands even though the original image is not available at the detector. Non-linear schemes treat both high and low magnitude coefficients with equal weight (only e.g. the sign of the coefficient is considered), thus suppression of image noise is achieved. Unlike linear detection methods using correlative processing (which would attach more significance to the high amplitude coefficients), in this case, large magnitude coefficients affect the result of the detection process the same way as the small magnitude coefficients.

3.4.1.3 Quantization index modulation methods

Some of the simplest watermarking algorithms, e.g. [249], which operate in the spatial domain and replace the least significant bits (LSB) of the image pixels, belong the same category that is discussed in the next sections. However, the quantize-and-replace strategies we will present below are much more advanced and allow for more robust watermarking.

As we have seen in the blind communication model of figure 3.15, the watermark message m is properly modulated and added to the host signal x. This embedding process can be written as the embedding function $s(x, m)$. However, we can also view $s(x, m)$ as a collection or ensemble of functions of x, indexed

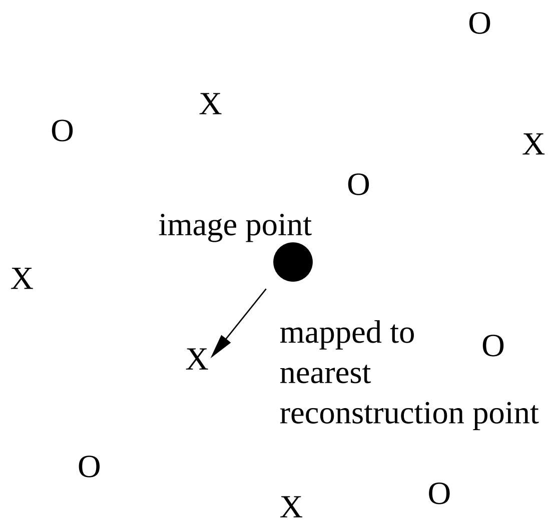

Table 3.3: Quantization index modulation. The reconstruction points are marked with x (for $m = 1$) and o (for $m = 2$) and belong to the two quantizers.

by m. To emphasize this view, we denote this ensemble of functions as $s(x; m)$. [28]

Due to the distortion constraint and imperceptibility requirement in watermarking applications, $s(x; m)$ should be close, at least perceptibly, to x for all m (approximate-identity property). The robustness requirement suggests, that the points in the range of one function in the ensemble should be "far away" in some sense from the points in the range of any other function. At least the ranges should be non-intersecting, otherwise it is not possible to determine the value of m from s.

Quantizers, or a sequence of quantizers, can be used to as approximate-identity functions to embed the watermark information. The number of possible values of m determines the number of required quantizers. m acts as an index that selects the quantizer that is used to represent m. For the case $m = 2$ (which we will discuss throughout the next sections) we have a binary quantizer.

Figure 3.3 illustrates the QIM information embedding process. To embed one bit m, $m \in \{1, 2\}$, an image pixel is mapped to the nearest reconstruction point representing the information of m.

The minimum distance d_{min} between the sets of reconstruction points of different quantizers in the ensemble determines the robustness of the embedding,

$$d_{min} = \min_{(i,j):i \neq j} \min_{(x_i, x_j)} \|s(x_i; i) - s(x_j; j)\|.$$

Intuitively, the minimum distance measures the amount of noise that can be tolerated by the system.

3.4.1.4 Dithered Quantization

Chen [28, 29] presented dithered quantization (or dither modulation) as a special case of quantization index modulation (QIM) for self-noise suppression. Dithered quantizers are quantizer ensembles where the quantization cells and reconstruction points of every quantizer in the ensemble are shifted versions of some base quantizer Q. The shift is given by a dither vector d.

Dithered quantization is an operation in which a dither vector d of length L is added to the input x prior to quantization. The output of the subtractive quantization operation is denoted by

$$s_i = Q(x_i + d_i) - d_i,\ 0 \leq i < L,$$

or, using the notation introduced above,

$$s(x; m) = Q(x + d(m)) - d(m).$$

For our discussion, we only consider uniform, scalar quantizer with a step size Δ. The binary dither ensemble can be generated pseudo-randomly by choosing $d_i(1)$ with a uniform distribution over $[-\Delta/2, \Delta/2]$ and assigning $d_i(2)$ as follows:

$$d_i(2) = \begin{cases} d_i(1) + \Delta/2 & d_i(1) < 0 \\ d_i(1) - \Delta/2 & d_i(1) \geq 0 \end{cases},\ 0 \leq i < L.$$

According to Schuchman [210], the subtractive dither quantization error (SDQE) does not depend on the quantizer input when the dither signal d has a uniform distribution within the range of one quantization bin ($d_i \in [-\Delta/2, \Delta/2]$, leading to an expected squared error $e^2 = \Delta^2/12$.

Algorithm 5 illustrates the dithered quantization of decoding operations, as well as the dither generation method outlined above.

3.4.1.5 Spread-transform dither modulation

The spread-transform dither modulation approach can be used to convert an existing spread-spectrum watermarking scheme into a scheme based on a quantize-and-replace strategy. This is achieved by simply replacing the addition with a quantization operation. Spread spectrum systems have an embedding formula of the form

$$s(x, m) = x + a(m) \cdot u,$$

where u is a normalized, pseudo-random vector.

The above embedding formula can be re-written in the form

$$s(x, m) = (\tilde{x} + a(m)) \cdot u + (x - \tilde{x} \cdot u),$$

where $\tilde{x}$ is the projection of the image x onto the spreading vector u, $\tilde{x} = x \cdot u$. Now, the addition step

$$\tilde{s} = \tilde{x} + a(m)$$

Algorithm 5 Chen's dither modulation algorithm.

```
double quantize(double value, double delta) {
  int q = rint(value / delta);
  return ((value - delta * q) <= (delta * (q + 1) - value)) ?
    delta * q :  delta * (q + 1);
}

void dm_quantize_vector(double x[], double dither[], double delta, int n)
{
  for (int i = 0; i < n; i++)
    x[i] = quantize(x[i] + dither[i], delta) - dither[i];
}

double dm_distance(double y[], double dither[], double delta, int n) {
  double sum = 0.0;
  for (int i = 0; i < n; i++)
    sum += sqr(y[i] - (quantize(y[i] + dither[i], delta) - dither[i]));
  return sum;
}

int decode_vector(double y[], double **dither, double delta, int n) {
  return (distance(y, dither[0], delta, n) <
    distance(y, dither[1], delta, n)) ?  0 :  1;
}

void generate_dither(double **dither, double delta, int n) {
  for (int i = 0; i < n; i++) {
    dither[0][i] = ranf() * delta - (delta / 2.0);
    dither[1][i] = dither[0][i] + (dither[0][i] < 0.0) ?
      (delta / 2.0) :  (-delta / 2.0);
  }
}
```

can be replaced by the quantization step

$$\tilde{s} = Q(\tilde{x} + d(m)) - dm$$

to convert the spread spectrum system to a STDM system. The final embedding formula of the STDM scheme is then

$$s(x; m) = (Q(\tilde{x} + d(m)) - d(m)) \cdot u + (x - \tilde{x} \cdot u).$$

Algorithm 6 implements the above quantization and decoding operations.

Algorithm 6 Chen's spread-transform dither modulation algorithm.

```
void stdm_quantize_vector(double x[], double u[], double dither, double
delta, int n) {
  double xp = proj(x, u, n);
  for (int i = 0; i < n; i++)
    x[i] = (quantize(xp + dither, delta) - dither) * u[i] +
      (x[i] - xp * u[i]);
}

double stdm_distance(double y[], double u[], double dither, double delta,
int n) {
  double yp = proj(y, u, n);
  return fabs(yp - (quantize(yp,delta) - dither));
}

void generate_spreading_vector(double u[], int n) {
  for (int i = 0; i < n; i++)
    u[i] = ranf() - 0.5;
  normalize(u, n);
}
```

3.4.1.6 Example: Algorithm Koch

Authors This algorithm has been developed by Eckhard Koch and Jian Zhao at the Fraunhofer Institute for Computer Graphics, Darmstadt, Germany and is published in [110, 274].

Watermark The mark is a sequence of binary values, $w_i \in \{0, 1\}$.

Coefficient selection The proposed algorithm pseudo-randomly selects 8×8 DCT coefficient blocks. Within each block b_i, two coefficients from the mid-frequency range are again pseudo-randomly selected (see figure 3.16). In later extensions to the basic scheme, certain block or coefficient pair are rejected based on robustness and watermark transparency criteria [274].

Embedding First, each block is quantized using to the JPEG quantization matrix and a quantization factor Q. Then, let f_b denote an 8×8 DCT coefficient block and $f_b(m_1, n_1), f_b(m_2, n_2)$ are the selected coefficients within that block. The absolute difference between the selected coefficients is given by

$$\Delta_b = |f_b(m_1, n_1)| - |f_b(m_2, n_2)|.$$

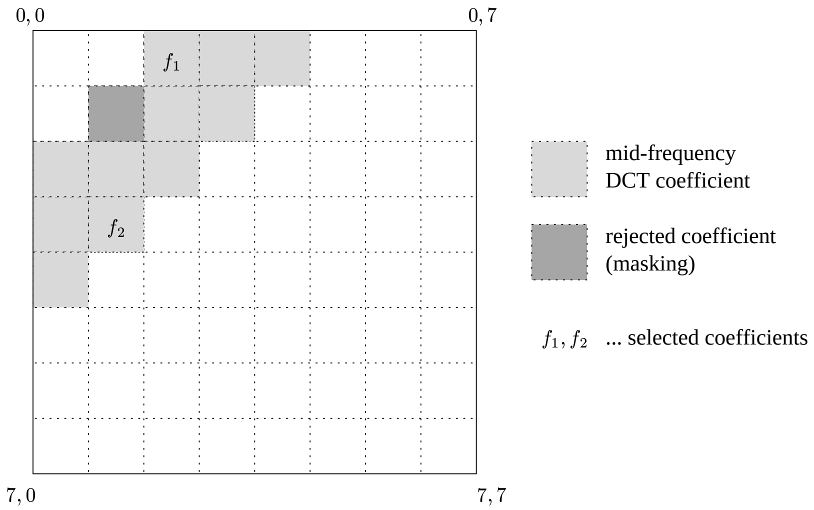

Figure 3.16: The basic watermarking scheme by Koch operates on 8×8 DCT coefficient blocks and manipulates a pair of coefficients to embed a single bit of watermark information.

In order to embed one bit of watermark information, w_i, in the selected block b_i, the coefficient pair $f_b(m_1, n_1), f_b(m_2, n_2)$ is modified such that the distance becomes

$$\Delta_b = \begin{cases} \geq q & \text{if } w_i = 1 \\ \leq -q & \text{if } w_i = 0 \end{cases},$$

where q is a parameter controlling the embedding strength.

Discussion The scheme has been extended by Benham [12] and Zhao [274] to enforce a relationship between three instead of two coefficients (see table 3.4). This improvement allows to encode the watermark bit in a more robust way and provides a technique to skip blocks that are not suitable for watermark embedding.
Since the watermark is embedded in 8×8 DCT domain coefficient blocks, visible artefacts may occur, especially in smooth regions.Further enhancements, e.g. [53], incorporate aspects of the HVS to tackle the above visible distortion problem. Very simple metrics have been proposed to capture the smoothness and edge activity of a block. Hence, a block can be either rejected or the embedding strength parameter q has to be adjusted.

3.4.2 Scalar Quantization

3.4.2.1 Algorithm Chu

Authors This algorithm has been developed by Chee-Jung Chu and Anthony Wayne Wiltz at the Center for Advanced Computer Studies, University of Southwestern Louisiana, Lafayette, LA, USA and is published in [31].

w_i	$f(m_1,n_1)$	$f(m_2,n_2)$	$f(m_3,n_3)$
1	H	M	L
1	M	H	L
1	H	H	L
0	M	L	H
0	L	M	H
0	L	L	H
X	H	L	M
X	L	H	M
X	M	M	M

Table 3.4: An extention to Koch's algorithm modifies the relationship between three coefficients. The information X is used to skip a block. L, M, H denotes the magnitude of the coefficients: low, middle and high, respectively.

(a) watermarking image (b) difference image

Parameter	Description
$N = 184$	length of the binary signature
$q = 5.0$	embedding strength
$Q = 10$	(pre-)quantization factor

Figure 3.17: Watermarked image (a) and difference image (b), created with Koch's algorithm and the embedding parameters above.

Watermark The mark is a sequence of binary values, $w_i \in \{-1, 1\}$.

Decomposition The authors propose a five-level integer wavelet decomposition.

Coefficient selection The algorithm used coefficient skipping for security and transparency reasons. Coefficients of the detail subbands are selected pseudo-randomly according to a density parameter. Only coefficients of the blue image component are manipulated because the human eye is least sensitive to blue color information.

Embedding The binary representation of the transform domain integer coefficients is right-shifted prior to embedding and left-shifted after the coefficient has been manipulated. This operation factors out the quantization distortion that the watermarked image is likely to undergo.
Each selected coefficient is modulated according to the following formula:

$$f'(m,n) = f(m,n) + \alpha \cdot l(m,n) \cdot w_i,$$

where α is a parameter determining the embedding strength and $l(m,n)$ is the luminance component of the host image at that location used to weight the watermark manipulation. The red-, green- and blue-channel information denoted by $r(m,n)$, $g(m,n)$, $b(m,n)$ is used to compute the luminance of a pixel, $l(m,n) = 0.299 \cdot r(m,n) + 0.587 \cdot g(m,n) + 0.114 \cdot b(m,n)$.

Extraction The watermark can be extracted without referring to an original image because the original blue channel coefficient can be estimated by averaging the coefficients in the neighborhood. The watermark information w_i is recovered by analyzing the sign of the difference between estimated, $\tilde{f}(m,n)$, and received coefficient, $f^\star(m,n)$,

$$w_i = \begin{cases} -1 & \tilde{f}(m,n) - f^\star(m,n) > 0 \\ 1 & \tilde{f}(m,n) - f^\star(m,n) \leq 0 \end{cases} .$$

Discussion The proposed scheme is very similar to a previous approach by Kutter [125] in the spatial domain.

3.4.2.2 Algorithm Hsu

Authors This algorithm has been developed by Chiou-Ting Hsu and Ja-Ling Wu at the National Taiwan University, Taiwan and is published in [81].

Watermark The watermark is a meaningful binary image such as a logo or a hand-written signature. The dimensions of the watermark image are assumed to be half of that of the host image.

Decomposition Both the watermark and the host image are decomposed successively into a multi-resolution representation. However, while a Daubechies-6 wavelet transform is used to decompose the host image, the binary logo image is decomposed with the resolution-reduction (RR) function of the

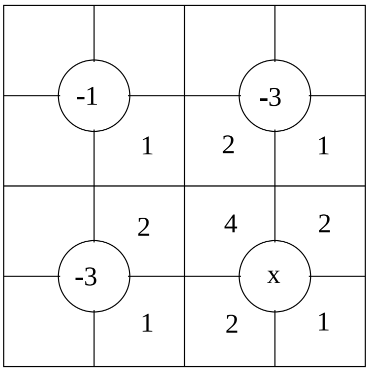

(a) JBIG resolution reduction

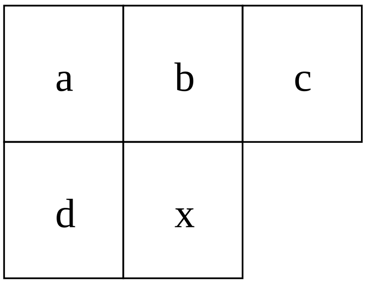

(b) residual mask

Figure 3.18: The JBIG resolution reduction technique (a) sets the target low resolution pixel x if the surrounding high- and low-resolution pixels have weighted sum greater than 4.5. High-resolution pixels are shown as circles, low-resolution pixels are shown as squares. (b) The residual mask used for embedding.

joint binary image experts group (JBIG) compression standard[3].This RR technique (depicted in figure 3.18 (a)) is more appropriate for bi-level images such as text or line drawings than intuitive sub-sampling by a factor of two because it preserves thin lines and other details.
After the resolution-reduction step, an up-scaled version of the residual is subtracted from the original watermark pattern in order to obtain the differential layer. Residual and differential layer will be used in the embedding stage.

Coefficient selection The differential layer and the residual of the watermark are embedded into the detail subbands of the host image at the same resolution. Note that the even columns of the watermark components are embedded into the HL_i subbands while the odd columns are cast onto the LH_i subbands. The host's approximation image and the HH_i subbands are not altered to avoid visible image distortion in the first case and due to the low robustness in the later case.

Embedding Before embedding, a pseudo random permutation is performed on the resolution-reduced version (the residual) and on the differential layer of the watermark image to disperse the spatial relationship of the binary pattern. Thus, a noise-like, statistically undetectable binary pattern is created.
The residual mask shown in figure 3.18 (b) is used to modify the neighboring relationship of host image coefficients. First, the residual polarity is computed between the neighboring pixels according to the residual mask. Then, the current coefficient (represented by the x in the mask) is changed to represent the corresponding watermark bit.

[3]ITU T.82

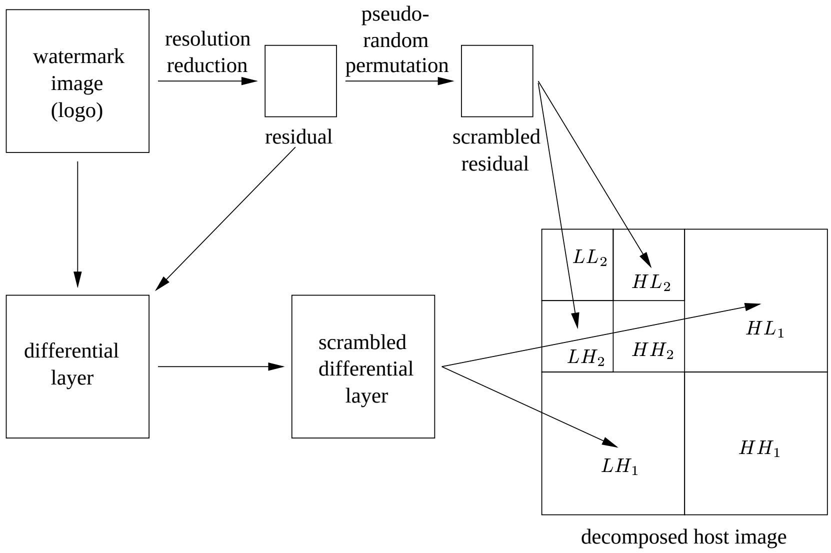

Figure 3.19: Hsu's scheme for binary watermark embedding.

Extraction Both the original and the watermarked image is required for watermark extraction. The proposed scheme is non-invertible without the knowledge of the host.

Discussion The authors point out that a meaningful logo image can not only be detected (using a correlative measure and a detection threshold) but also extracted.
The resolution-reduced watermark image is embedded into the LL subband of the host image. In the detail subbands of the host, the differential layer of the watermark pattern is hidden. This ways, the subband characteristics of host and watermark image are matched, resulting in imperceptible information embedding.
The choice of wavelet filters is said to affect the quality of the watermarked image and the robustness to attack. In image compression systems, filters which pack most of the image's energy in the approximation image, are preferred. Using such filters in the above watermarking method will seriously harm its robustness because the watermark information is embedded in the detail subband only. Therefore, a compression system can easily discard the perceptually irrelevant watermark information.

3.4.2.3 Algorithm Inoue

Authors This algorithm has been developed by Hisashi Inoue, Akio Miyazaki, Akihiro Yamamoto and Takashi Katsura at the Kyushu Multimedia System Research Laboratory, Matsushita Electric Industrial Company, Iizuka, Japan and is published in [87, 84, 85, 88, 86, 83].

Watermark The mark is a sequence w_i of binary values matching the number of zerotrees found in the host image.

Decomposition The authors propose a 3-level decomposition using 5/3 symmetric short kernel filters (SSKF) or Daubechies-16 filters.

Coefficient selection Wavelet coefficients are classified as significant or insignificant using the notion of zerotrees. The zerotree structure has been first described by Lewis [132] and successfully incorporated in image compression schemes by Shapiro [212] (see section 2.7.3.2).
Given the threshold T, a wavelet coefficient $f(m,n)$ is said to be insignificant if $|f(m,n)| < T$. If a coefficient and all of its descendants (i.e. the coefficients corresponding to the same spatial location but at a finer scale of the same orientation) are insignificant with respect to T, then the set of these insignificant wavelet coefficients is called zerotree for the threshold T.
A coefficient at the coarse scale is called parent while all four coefficients at the finer resolution level are called children. A wavelet coefficient is called zerotree root if it is not the descendant of a previously found zerotree root.
Let f_{max} denotes the maximum absolute wavelet coefficient value of a set of subbands, needed to adjust the significance threshold $T = \alpha \cdot f_{max}$ to a given image, where α is a scaling factor, $0.01 < \alpha < 0.1$.
The proposed watermarking scheme comes in two flavors, one based on manipulating significant (method A) and the other one based on manipulating insignificant coefficients (method B).
For method A, all zerotrees Z_i for the threshold T are selected, not taking the lowest frequency subband (LL) into account. Method B selects significant coefficients from the coarsest scale detail subbands (LH_3, HL_3, HH_3). The coefficients are selected such that $T_1 < |f(m,n)| < T_2$, where $T_2 > T_1 > T$.

Embedding The embedding strategy for method A sets all coefficients of zerotree Z_i to $-m$ to signal $w_i = 0$ and sets the coefficients of Z_i to m if w_i is 1.
Method B casts the watermark symbol w_i on a selected coefficient via quantization according to the following rule.

$$f'(m,n) = \begin{cases} T_2 & w_i = 1 \, and \, f(m,n) > 0 \\ T_1 & w_i = 0 \, and \, f(m,n) > 0 \\ -T_2 & w_i = 1 \, and \, f(m,n) < 0 \\ -T_1 & w_i = 0 \, and \, f(m,n) < 0 \end{cases}$$

Extraction The watermark extraction algorithms of method A computes the average coefficient value M_i for the coefficients belonging to zerotree Z_i.

$$w_i = \begin{cases} 0 & M_i < 0 \\ 1 & M_i \geq 0 \end{cases}$$

Method B extracts the watermark symbol w_i from a significant coefficient $f^\star(m,n)$ by checking the magnitude.

$$w_i = \begin{cases} 0 & |f^\star(m,n)| < (T_1 + T_2)/2 \\ 1 & |f^\star(m,n)| \geq (T_1 + T_2)/2 \end{cases}$$

(a) watermarked image

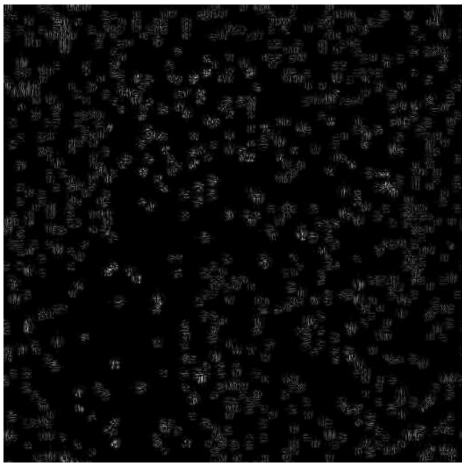

(b) difference image

Parameter	Description
$\alpha = 0.05$	significance threshold scaling
$M = 6.0$	embedding strength

Figure 3.20: Watermarked image (a) and difference image (b), created with Inoue's algorithm (semi-blind, insignificant coefficients, method A) and the embedding parameters above.

Discussion Although the authors claim that method A of the algorithm can extract the embedded information without the original host image (blind extraction), the scheme uses the positions of zerotree roots to guide the extraction algorithms. Therefore, the algorithm should be actually called semi-blind. Without this guidance, the proposed algorithm looses synchronization easily because it depends on insignificant coefficients only.
This shortcoming can be relatively simple removed by enforcing additional rules before accepting a zerotree location. In particular, we reject a zerotree if the magnitude of at least one of its coefficients is larger than $\gamma_1 \cdot M$, if the sum of all coefficients in the zerotree is less than $\gamma_2 \cdot M$ or if the difference between the minimum and maximum coefficient in the zerotree set is larger $\gamma_3 \cdot M$. Experimentally, we found that satisfactory blind watermark extraction results can be obtained by setting $\gamma_1, \gamma_2, \gamma_3$ to $3.5, 0.1$ and 3.5, respectively.

3.4.2.4 Algorithm Kundur

Authors This algorithm has been developed by Deepa Kundur and Dimitrios Hatzinakos at the Department of Electrical and Computer Engineering, University of Toronto, ON, Canada and is published in [115, 116, 119, 112, 111].

Watermark The watermark is a sequence of binary values.

(a) watermarked image

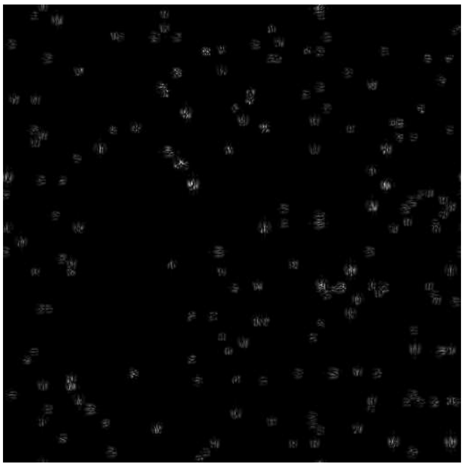

(b) difference image

Parameter	Description
$\alpha = 0.1$	significance threshold scaling
$M = 6.0$	embedding strength

Figure 3.21: Watermarked image (a) and difference image (b), created with Inoue's algorithm (blind, insignificant coefficients, method A) and the embedding parameters above.

Decomposition The authors propose using the Daubechies family of orthogonal wavelet filters to derive a multi-resolution representation of the image data. A decomposition level of 3 or 4 is used in the experiments.

Coefficient selection The algorithm pseudo-randomly selects locations in the detail subbands. Each time, a coefficient triple $(f_{h,l}(m,n),\ f_{v,l}(m,n),\ f_{d,l}(m,n))$ is selected from three distinct detail subbands of one decomposition level. See figure 3.23.

Embedding The selected coefficient triple is sorted in ascending coefficient magnitude order. Then the median coefficient is quantized to represent the information of a single watermark bit, w_i. As illustrated in figure 3.24, the median coefficient is set to the nearest reconstruction point that represents the current watermark information.
The bin width parameter Δ controls the quantization step size. Coarser quantization will lead to more robust watermark embedding, however, this will also introduce more distortion.

Extraction The proposed algorithm features blind watermark extraction.

Discussion In order to improve robustness, an extension of the above scheme is described in [116, 119, 111]. A known reference watermark is embedded together with the secret watermark in an interlaced way. It is assumed that both watermarks, i.e. the reference and the secret mark, undergo the same type and amount of distortion due to attack or image processing. Thus, by

(a) watermarked image

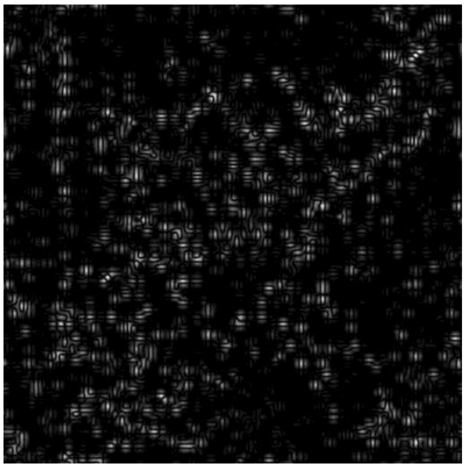

(b) difference image

Parameter	Description
$\alpha = 0.1$	significance threshold scaling
$T_1 = 40$	lower quantization boundary
$T_2 = 90$	upper quantization boundary

Figure 3.22: Watermarked image (a) and difference image (b), created with Inoue's algorithm (significant coefficients, method B) and the embedding parameters above.

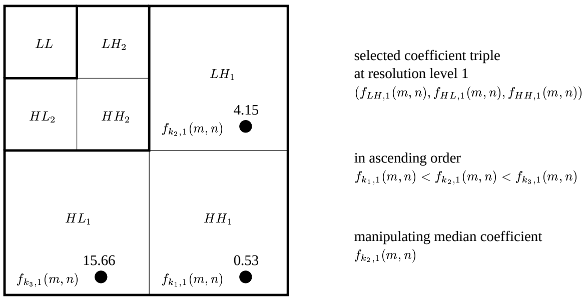

Figure 3.23: Coefficient selection in the watermarking scheme by Kundur.

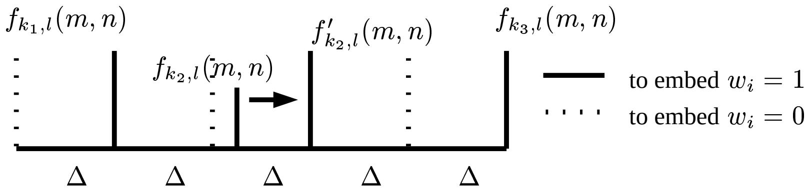

Figure 3.24: The bin quantization technique of Kundur's watermarking scheme.

analyzing the error of the known reference watermark, one can estimate the error on the secret watermark. This allows to apply different weights to the recovered watermark information. Furthermore, the watermark is redundantly embedded to improve robustness.
In [121], the authors claim that by embedding the watermark in a different domain than the one used for image compression, the robustness can be improved – at least the watermarking and compression systems should use different perceptual models.

3.4.2.5 Algorithm Kundur (fragile)

Authors This algorithm has been developed by Deepa Kundur and Dimitrios Hatzinakos at the Department of Electrical and Computer Engineering, University of Toronto, Canada and is published in [118, 117, 120, 112].

Watermark A binary signature is embedded into key-selected detail subband coefficients.

Decomposition The authors obtain a multi-resolution representation of the host image using an integer Haar wavelet transform.

Coefficient selection A pseudo-random coefficient selection algorithm is used, leaving some coefficients unmodified to limit visual distortion of the image.

Embedding The schemes builds on the quantization method [115] described in section 3.4.2.4. An integer wavelet transform is employed to avoid round-off errors during the inverse transform which could be detected as an tampering attempt.

(a) watermarked image (b) difference image

Parameter	Description
$N = 184$	watermark length
$L = 2$	decomposition level
$Q = 3$	quantization factor

Figure 3.25: Watermarked image (a) and difference image (b), created with Kundur's algorithm and the embedding parameters above.

Discussion The proposed algorithm is designed to be fragile. Due to the spatial localization of the DWT coefficients, this tamper-proofing scheme can detect the spatial as well as the frequency regions that have been tampered with.

3.4.2.6 Algorithm Matsui

Authors This algorithm has been developed by Kineo Matsui, Junji Ohnishi and Yasuhiro Nakamura at the Department of Computer Science, National Defense Academy, Yokosuka, Japan and is published in [156].

Watermark The mark is a sequence of bits, $w_i = \{0, 1\}$.

Decomposition The host image is Haar wavelet filtered to a multi-resolution representation.

Coefficient selection The algorithm selects coefficients $f_{LH}(m, n)$, $f_{HH}(m, n)$, $f_{HL}(m, n)$ from the detail subbands at the same resolution scale and forms a vector $V_i = (f_{LH}(m, n), f_{HH}(m, n), f_{HL}(m, n))$. According to table 3.5, a class $C(V_i)$ can be associated to each vector V_i.

Embedding Beforehand, a specific vector class c has to be chosen for embedding, except class 0 which can not be used for embedding purposes. For

class	LH	HH	HL
0	0	0	0
1	0	0	X
2	0	X	0
3	0	X	X
4	X	0	0
5	X	0	X
6	X	X	0
7	X	X	X

Table 3.5: Classification table for detail subband vectors $(f_{LH}(m,n), f_{HH}(m,n), f_{HL}(m,n))$ according to Matsui's scheme. X denotes a non-zero coefficient value.

each V_i the corresponding vector class $C(V_i)$ is determined. If $C(V_i)$ does not match c than this coefficient has to be skipped. Otherwise, the first non-zero element, say v, in the vector V_i is modified according to the following rules. The lowest k bits of the coefficient v are set to w_i. If v is now zero then the $(k+1)$th bit is set to 1. If there are other non-zero elements in V_i manipulate them accordingly.

Extraction The inverse procedure is used for watermark extraction.

Discussion The classification scheme can be used to distinguish between natural and artificial (computer generated) images since artificial images have a preponderance of coefficient vectors in class 0.
The k least-significant bits of each selected non-zero vector element are replaced with one signature bit. Therefore, the proposed algorithm allows good capacity and low perceptual distortion. On the other hand, it is not very robust against attacks.

3.4.2.7 Algorithm Ohnishi

Authors This algorithm has been developed by Junji Ohnishi and Kineo Matsui at the Department of Computer Science, National Defense Academy, Yokosuka, Japan and is published in [170, 171].

Watermark The watermark is sequence of bits, $w_i \in \{0,1\}$.

Decomposition A two-level multi-resolution representation of the host image is obtained using a Haar wavelet transform.

Coefficient selection All coefficients of the approximation image (LL subband) are selected and subjected to a pseudo-random noise sequence n, where $n_i \in \{-1,1\}$. $\tilde{f}(m,n) = f(m,n) \cdot n_i$. The product of $\tilde{f}(m,n)$ with the same noise sequence n_i reconstructs the original signal because

$$\tilde{f}(m,n) \cdot n_i = f(m,n) \cdot n_i^2 = f(m,n)$$

since $n_i \in \{-1,1\}$.
Next, the spread transform coefficients $\tilde{f}(m,n)$ are segmented into blocks

of dimension $B \times B$. Each block b_j is manipulated separately to carry one watermark bit, w_i.

Embedding The Fourier transform is applied individually on each block to compute its frequency representation $\tilde{b}_j(k,l)$. To embed a single watermark bit, the DC coefficient of a block $\tilde{b}_j(k,l)$ is uniformly quantized with step size Δ.

Extraction The inverse procedure is used to extract the watermark sequence (without the need of the original image).

3.4.2.8 Algorithm Xie

Authors This algorithm has been developed by Liehua Xie and Gonzalo R. Arce at the Department of Electrical and Computer Engineering, University of Delaware, Newark, DE, USA and is published in [268, 267].

Watermark The watermark is a sequence of binary values, $w_i \in \{0, 1\}$.

Decomposition The host image is decomposed to obtain a low-frequency approximation representation.

Coefficient selection The watermark is embedded solely in the approximation image (LL subband) of the host image. Each time, the coefficient triple of a non-overlapping 3×1 sliding window is selected and manipulated. This process is shown in figure 3.26.

Embedding First, the elements b_1, b_2, b_3 of the local sliding window are sorted in ascending order according to their magnitude. Then the range between $\min|b_j|$ and $\max|b_j|$, $j = 1 \ldots 3$ is split into intervals of length

$$\Delta = \alpha \cdot \frac{\max|b_j| - \min|b_j|}{2}.$$

Next, the median of the coefficient triple is quantized to become a multiple of Δ in order to represent one bit of watermark information, w_i. Thus, the interval is divided into $\frac{2}{\alpha}$ regions, where each has region has two boundaries, l_k and l_{k+1}. Now we associate the one bit to all even boundaries and the zero bit to all odd boundaries of the interval. The boundaries are also named reconstruction points since they are part of the output set of the quantizer. The median coefficient is modified to lie on a boundary representing the information of watermark bit w_i. Finally, the manipulated coefficient is updated in the host image's subband.

Extraction The watermark extraction algorithm works without the original image. The median of the sliding window is determined and quantized to obtain a reconstruction point. The bit value associated with that reconstructed point is extracted and assigned to w_i.

Discussion Although the schemes is proposed for image authentication application, however, contrary to the authors' claims, its robustness makes it also suitable for other purposes such as copyright protection. We found that the robustness is mainly determined by the number of decomposition

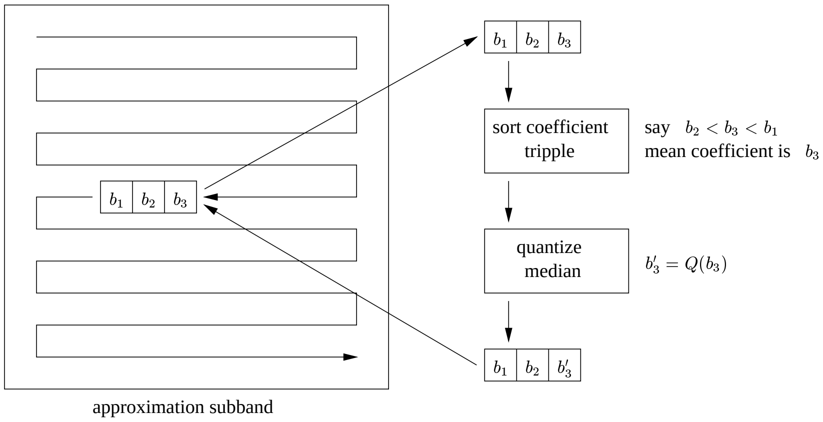

Figure 3.26: Watermarking scheme by Xie, showing the bit "engraving" method.

steps . Very good robustness can be achieved employing e.g. a five-level wavelet decomposition using Daubechies-7/9 bi-orthogonal filters. Embedding at two resolution levels, level 2 and level 5, is illustrated in figure 3.27 and 3.28, respectively.
The authors discuss employing the above algorithm in the EZW [267] and SPIHT [268] coding schemes (these compression schemes have been briefly reviewed in section 2.7.3).
Not only binary embedding, but also m-ary signaling is proposed and compared with an analytical capacity bound. The authors found that multi-bit "engraving" can improve the capacity.

3.4.3 Vector Quantization

In the last section, the quantizer inputs were scalar values and each quantizer codeword represented a single sample of the source output. A quantization strategy that works with sequences or blocks of output is called vector quantization. The problem is to generate a representative set of sequences, called codebook. Given a source sequence or source vector, we would represent it with one of the elements in the codebook. See figure 3.29 for an illustration of that process.

The quantization algorithm has to find the closest vector in the codebook for a given source vector, which can be computationally expensive if the codebook is large. To facilitate the search, the codebook is usually structured in some way. In the following, lattice points are used as a structure for vector quantization.

3.4.3.1 Algorithm Chae

Authors This algorithm has been developed by Jong Jin Chae, Debargha Mukherjee and B. S. Manjunath at the Department of Electrical and Com-

(a) watermarked image (b) difference image

Parameter	Description
$\alpha = 0.5$	quantization step size
$L = 5$	decomposition steps
$N = 80$	watermark length

Figure 3.27: Watermarked image (a) and difference image (b), created with Xie's algorithm (robust "etching") and the embedding parameters above.

puter Engineering, University of California, Santa Barbara, CA, USA and is published in [25, 24, 26, 22, 23, 21].

Watermark The mark is a sequence w_i of β-ary symbols. The watermark sequence is derived from a logo image, a quarter of the size of the host image. A one-level DWT decomposition of the logo image is computed and the coefficients are quantized into β levels.

Decomposition The authors propose using a one-level decomposition with the Haar wavelet filter.

Coefficient selection n transform coefficients are grouped together to form a n-dimensional vector. In particular, the case $n = 4$ is discussed in the paper, resulting in the D_4 lattice structure. To embed the quantized information of one logo image coefficient, one host image vector is manipulated. The coefficients of the logo image's LL subband are embedded by perturbing vectors from the LL subband of the host image. The information of the detail subbands is hidden in the corresponding detail subbands of the host image with the same orientation.

Embedding A vector of DWT host coefficients, v, is modified according to the scaled codeword representing w_i,

$$v' = v + \alpha \cdot C(w_i).$$

(a) watermarked image (b) difference image

Parameter	Description
$\alpha = 0.2$	quantization step size
$L = 2$	decomposition steps
$N = 5380$	watermark length

Figure 3.28: (a) Watermarked image and (b) difference image, created with Xie's algorithm (fragile "engraving") and the embedding parameters above.

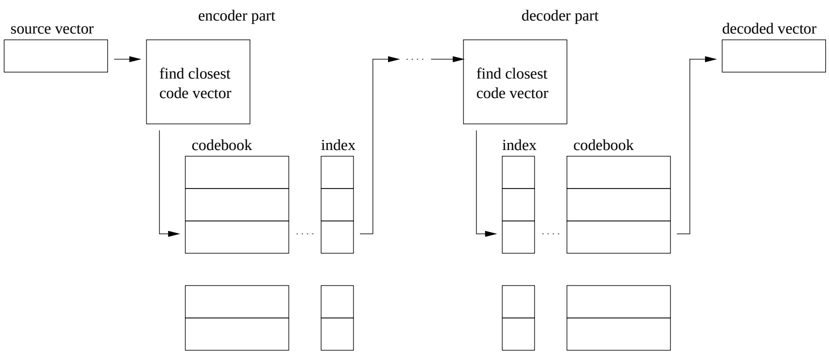

Figure 3.29: The vector quantization procedure.

Thus, for $n = 4$, four coefficients of the host image are modified to embed one quantized coefficient of the logo image.

Extraction The scheme has to have access to the original image in order to extract the embedded logo image. The error vector

$$e = \frac{v^\star - v}{\alpha}$$

is computed and used in a nearest-neighbor search against the codebook to reconstruct the embedded information

$$w_i = min_{w_i} \|C(w_i) - e\| .$$

Discussion The proposed scheme, especially the coefficient selection method (see figure 3.12), is similar to the algorithm by the same author described in section 3.3.3.1. However, the vector quantization approach is much more flexible and allows to control the robustness or distortion level and the quality of the embedded logo image via parameter α (embedding strength) and parameter β (quantization level), respectively.

3.4.4 Miscellaneous Algorithms

3.4.4.1 Algorithm Lin

Authors This algorithm for image authentication and tamper detection has been developed by Ching-Yung Lin and Shih-Fu Chang at the Department of Electrical Engineering, Columbia University, NY, USA and is published in [136].

Decomposition The schemes is based on a one-level wavelet decomposition.

Embedding The entire HH subband of the host image is replaced with a pseudo-random noise pattern. Image manipulation can be detected and localized by analyzing the integrity of the embedded pattern. The authors propose using 16×16 pseudo-noise blocks which are embedded repetively.

Discussion The scheme relies on a secret transform structure that secures the embedded noise pattern.

3.5 Discussion

3.5.1 General low-frequency subband algorithms

Several algorithms have been proposed that embed a watermark in the low-frequency subband or multi-resolution approximation image of the host image. Transforming an image with a 8×8 DCT can be seen to produce hierarchical data equivalently to a three level subband transform of 64 frequency bands [240]. The DC coefficients of all transformed blocks of image data represents the low-resolution approximation image. Figure 3.30 illustrates that process. Thus, a

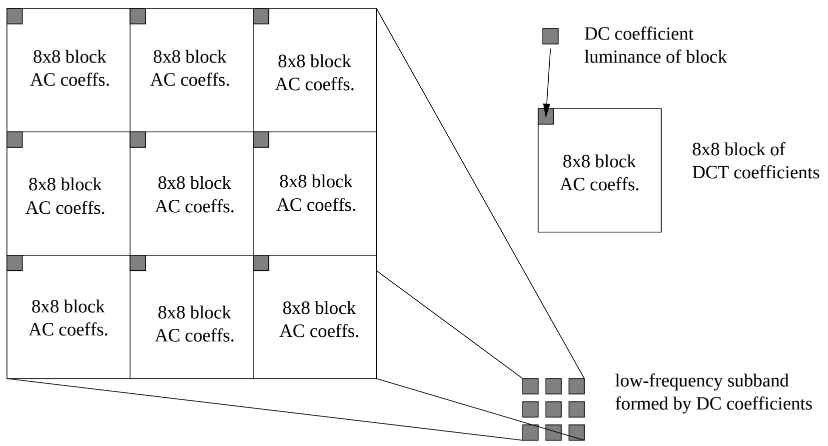

Figure 3.30: The approximation "subband" formed of the DC coefficients of several 8×8 blocks of DCT domain coefficients.

DCT on 8×8 image blocks can be used as an alternative to a three-level wavelet decomposition if only the LL subband is to be watermarked.

One of the first transform domain watermarking approaches taking the wavelet representation into consideration has been proposed by Ruanaidh [174] and Bowland [15]. Their algorithm operates on non-overlapping image blocks and embeds a sequence of binary symbols by means of bi- or uni-directional coding [15, 182].

Liang et al. [134] proposed a watermarking scheme that can be used to embed a watermark in the low-frequency component of the host image. The low-frequency component can be accessed in the several transform domains obtained from e.g. the DCT, a lapped transform (the author uses the 16×8 lapped bi-orthogonal transform, LBT), or the DWT (here several resolutions are easily possible, depending on the number of decomposition steps). Regardless of the actual transform involved, the low-frequency subband is subjected to a DFT and the watermark is linearly embedded in the magnitude coefficients of the DFT transform. Results indicate that watermarking in the DWT approximation subband is superior to watermarking other low-frequency representations of the host image.

3.5.1.1 Algorithm Pereira

Authors This algorithms has been developed by Shelby Pereira, Sviatoslav Voloshynovskiy and Thierry Pun at the University of Geneva, Switzerland and is published in [177].

Watermark The watermark is a sequence of bits.

Decomposition The scheme depends on a one-level decomposition of non-overlapping 16×16 image blocks using Haar wavelet filters.

Coefficient selection For each bit to be embedded, a 2×2 block of neighboring coefficient is selected from a LL subband which has size 8×8. 64 coefficients are available to form our 2×2 blocks, however, only 8 embedding blocks are used. It is important to select neighboring coefficients since it is assumed that their difference is 0 on average.

Embedding A watermark bit is placed in a 2×2 embedding block using differential encoding.

Discussion The proposed watermarking algorithm uses linear programming to optimize watermark robustness within a visual distortion constraint given by JND threshold maps.

3.5.1.2 Algorithm Tzovaras

Authors This algorithm has been developed by D. Tzovaras, N. Karagiannis and M. G. Strintzis at the Electrical and Computer Engineering Department, University of Thessaloniki, Greece and is published in [240].

Watermark The watermark is determined by the seed value for the pseudo-random number generator.

Decomposition The authors compare using a block-based 8×8 DCT and a three-level wavelet decomposition with Haar and bi-orthogonal filters. Figure 3.30 depicts how the approximation subband is derived using the DCT.

Coefficient selection All coefficients from the low-resolution approximation image are selected and split into two subsets, A and B, according to a pseudo-random sequence of binary values. Furthermore, only coefficients are selected which belong to regions in the image that contain sufficient texture information, i.e. blocks whose AC energy is greater than a specific threshold.

Embedding The embedding process is similar to the approach described by Pitas [182]. The coefficient values of one subset are increased while the coefficients of the other subset are decreased in order to maximize the difference of the sample means of the two subsets.

Extraction Only the presence or absence of the mark can be detected without referring to the original image using hypothesis testing. Without an embedded watermark, the mean values of the two subsets are expected to be about the same. A mark is said to be detected if the mean value of subset A is significantly different from subband B.

Discussion The proposed watermarking method demonstrates that the DCT and DWT can both be employed when embedding the watermarking only in the low-resolution approximation image.

3.5.2 Perspective

A number of research work related to watermarking in the wavelet domain has been published only very recently. In this section, we try to briefly summarize these novel approaches.

Chen et al. [30] proposed a watermarking method, embedding a sequence of binary values in significant wavelet coefficients which have been determined by Shapiro's [212] zerotree algorithm. They discuss the advantages of transform domain watermarking algorithm in general and the benefits of the wavelet domain in particular. Furthermore, the security aspects of the scheme are analyzed, stating that the specific wavelet transform used in the decomposition step is a key element to its reliability.

Jayawardena et al. [89, 90] successively apply binary wavelet filters [226] to obtain a multi-resolution domain. The decomposition step follows the lifting approach [230] and incorporates the binary XOR operator to predict the difference between even and odd samples. The algorithm selects a significant bit-layer of the detail subbands. First, all bits in that layer are set to one and the inverse wavelet transform is applied. The resulting image is denoted by I_1. Next, all bits in the selected layer are set to zero and, again, the inverse transform is applied to compute I_0. Now, observe for which location of the selected bit-plane the embedded information is "stable" when the image is subjected to lossy image compression. Moreover, a given visual distortion bound has to be respected for each prospective bit storage location. The set of "stable" locations is used to directly embed the binary watermark.

Tsekeridou et al. [239] exploit the multi-resolution property of the wavelet transform domain and embed a circular self-similar watermark [214] in the first- and second-level detail subbands of a wavelet decomposition. The self-similarity proves useful for watermark detection without the original image since the search-space to locate the embedded watermark can be drastically reduced if the image has undergone geometric distortion.

The wavelet packet transform based on the best-basis algorithm [260] is used for a novel embedding method proposed by **Manoury** [152] and **Vehel** [131]. Here, the watermark information is embedded by manipulating the structure of the wavelet decomposition. The energy of certain subbands in a decomposition sub-tree is manipulated in order to coerce a decomposition structure which represents a given watermark bit.

Based on his scheme described in [237], **Tsai** exploits the adaptivity of the best-basis wavelet packet transform to improve security and robustness [238]. The algorithm embeds the logo-type watermark only in the most important and most complex region of an image in order to avoid making visible changes to the background. These regions are determined in the spatial domain. The wavelet packet transform is employed to capture the energy of the selected regions in the transform domain.

The local contrast of an image can be computed in the wavelet transform domain using analytic filters [241]. **Vandergheynst** demonstrated that the proposed contrast model can be used in existing watermarking schemes [122] to weight the embedding strength according to local image activity.

Davoine [49] compared a watermarking schemes similar to the one proposed by Kundur [115] to a new technique which partitions the union of the lowest-resolution detail subbands into distinct regions that have approximately the same number of significant coefficients. The significant coefficients of a region are quantized to represent one bit of watermark information. The later algorithm has the advantage of flexibility, because it is not limited to quantizing

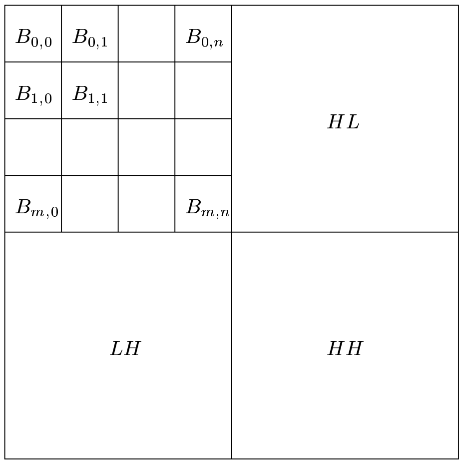

Figure 3.31: The wavelet packet decomposition used in Ejima's watermarking scheme.

coefficient triples but can adapt the number of significant coefficients per region to robustness requirements. However, both approaches only achieve semi-blind watermark extraction since reference data is required: the locations of the coefficient triples in the first case or the partitioning of the detail subbands in the second case.

A specific wavelet packet decomposition is used for image and video watermarking by **Ejima** [67, 65, 66] (see figure 3.31). The energy for each subband $B_{i,j}$ is computed and certain subbands are pseudo-randomly selected according to their energy. The mean absolute coefficient value of each selected subband is quantized and to encode one bit of watermark information. Finally, pseudo-randomly selected coefficients of that subband are manipulated to reflect the quantized coefficient mean value.

Inoue et al. [86] proposed both, a fragile and a robust watermarking scheme for tamper proofing based on their previous work [87, 85, 88, 83]. The robust approach quantizes coefficients of the multi-resolution approximation image to encode the binary watermark information. The fragile method selects significant coefficients above a given threshold T from the detail subbands (excluding the finest resolution subbands). These significant coefficient are then quantized the same way as for robust watermarking while insignificant coefficients are replaced by the product of the watermarking bit $w_i \in \{-1, 1\}$ and constant m below the threshold T. By analyzing the extracted watermark of the received image, image distortion can be classified in unintentional changes and malicious attacks.

A watermarking system that is integrated in the EBCOT [232] coding pipeline is proposed by **Su** [224]. Since the EBCOT coding scheme has be incorporated into JPEG2000, this method is directly applicable to the forthcoming compression standard. Basically, Su's scheme adds a Gaussian sequence to significant coefficient of a code-block. To fit into the JPEG2000 coding pipeline (shown in figure 2.2), each code-block has to be treated individually. Blind watermark detection is possible by computing the mutual correlation between the received watermarked coefficients and the watermark sequence itself. It is important to note that many significant coefficients have to be correlated in order to reliably detect the watermark after the host image was subjected to attacks. However, due to the hierarchical organization of the bitstream, it might be possible to detect the watermark solely from code-blocks belonging to the low-resolution subbands.

Kim [105] presented a multi-resolution watermarking method where the number of coefficients manipulated in each subband is proportional to the subband's energy. The total number of watermarked wavelet domain coefficients is set to 5000. Using a perceptually weighted additive embedding formula, 500 coefficients in the LL subband are watermarked. The remaining 4500 locations are allocated to the multi-resolution detail subbands according to their energy.

An image adaptive watermarking method is proposed by **Kaewkamnerd** [92] that embeds the mark, a pseudo-random Gaussian sequence, in the low-resolution detail subbands. A visual masking model [4] is used to weight the watermark embedding process in order to achieve imperceptibility although significant subbands are modified.

3.5.3 Further concepts

Kutter [124] suggested that "second generation" watermarking schemes should embed the watermark in significant data features. He proposed a method that extracts feature points from images which are invariant to common geometrical transformations, such as scaling, rotation, cropping, ... A multi-resolution analysis employing the rotation-invariant Mexican-hat wavelet detects feature point at various scales. Next, a Voronoi diagram based on the detected feature points is computed in order to segment the image. Each individual segment is then watermarked with a scheme similar to those proposed in [125, 123, 31]. This results in object-based watermarking, one of the main concepts of "second generation" schemes.

Kanai [99] used a wavelet transform ("lazy wavelets") to decompose the coefficient vectors of a 3D polygon model and embed a binary watermark in the resulting multi-resolution representation of the model.

Video watermarking based on a temporal wavelet transform is discussed in the work of Swanson [227, 228]. Here, a group of frames (a scene) is separated by the wavelet transform in static and dynamic components.

Piva describes an object-based watermarking system for MPEG-4 video streams built on the DWT [187]. The system allows detection of a watermarked object that has been copied to another stream. Since the scheme operates on a frame-by-frame basis, it is also applicable to image watermarking.

The following papers were not available for our research or have not been published yet: [67, 170, 203, 231, 263, 30]

Chapter 4

Contributions

In this chapter, two concepts are discussed which can improve the security, performance and applicability of watermarking schemes. Security issues of many, especially blind, watermarking schemes are addressed in section 4.2 where Fridrich's idea [71, 70] of key-dependent basis functions is presented. This work extends his concept to the wavelet domain and proposes to use a key-dependent wavelet transform in section 4.3.

In section 4.4, we investigate watermarking integrated in the image coding process and propose a novel watermarking scheme that is compatible with the operation of the upcoming JPEG2000 image coding standard. Two application scenarios, namely copyright protection and image authentication, are demonstrated and show the robustness and capacity of our embedding method.

4.1 Security Issues

In watermarking applications, the following operations raise security issues and and need to be protected to withstand an malicious attacker:

- watermark detection,
- watermark extraction,
- watermark modification and
- watermark removal.

It is important to note which information is needed to perform the above operation. This information can comprise knowledge about the particular watermarking algorithm, information about some designated key material, knowledge of the transform structure used to obtain the transform domain in which the watermark was embedded, knowledge about certain parameter used in the embedding process (such as embedding strength, quantization step size, ...) and knowledge of the coefficient selection algorithm. Furthermore, access to the original image or to other, similarly watermarked images, is also a security consideration. The technique of marking the same image with different watermarks (e.g. for customer tracking) is called fingerprinting and raises security problems of its own right because the fingerprint has to be collusion secure [16, 52, 181]. Especially in video watermarking applications, many similar frames are watermarked, opening the chance of a specialized attack.

A watermark attack is an operation on the watermarked image or an operation within the underlying protocol layer, that is aimed to break the intended purpose of the watermarking application. All distortion applied to the watermarked image has to be seen as an attack. Therefore, we have to distinguish between intentional or hostile and unintentional attacks. The sole purpose of hostile watermark attacks is to break the watermark while unintentional attacks might occur during normal image processing operations, such as compression. A more detailed classification of watermark attacks is presented in chapter 5.

When discussing weaknesses of watermarking algorithms, we have to differentiate between robustness problems and security problems.

The robustness is the property of a watermarking scheme to withstand image distortion. Ideally, a watermarking application for copy protection should be able to correctly extract or detect the watermark as long as the host image is still valuable and usable [179, 178]. If the image distortion grows to a level that renders the image useless then the watermark is allowed to fail. For image authentication applications this means that the watermark has to tolerate unintentional distortion which does not change the 'meaning' or the intelligible content of the image. As soon as certain regions of the image are 'forged', the watermark is intended to break and reveal the manipulated areas.

Security problems arise from weaknesses in the watermarking algorithm itself. An attacker who can exploit his knowledge about working principles of the watermarking algorithm can break the algorithm much more efficiently and with less distortion (compared to less sophisticated attacks against the robustness of the watermarking scheme).

The following section will introduce Fridrich's idea [71, 70] of key-dependent basis functions. His method tries to improve the security of watermarking applications with regards to the following problems:

Manipulation of specific coefficient locations. Blind watermarking algorithms often quantize coefficients at chosen locations in the transform domain to embed the watermark information. If these coefficient locations can be guessed or predicted, than the watermark can be easily destroyed or manipulated while adding only modest amounts of image distortion. Especially watermarking methods that weight the strength of the mark in an image-adaptive way are vulnerable to this kind of attack because the attacker must be assumed to have the same knowledge about the significance of certain coefficients as the embedding algorithm [244, 245, 242]. To counter this attack, most watermarking schemes employ pseudo-random coefficient skipping. However, this approach deceases the capacity of the watermark and also the robustness.

Image-adaptive threshold thwarting. Certain image-adaptive schemes select the regions (e.g. subbands) where the watermark is to be embedded according to a computed significance threshold. If the threshold computation is based on just a few coefficients (e.g. the maximum coefficient per subband, as for example in Wang's scheme [256]) and the original image is not available as a reference, the watermarking scheme can be attacked by thwarting the significance computation. In a later work, Wang improves the significance computation [254].

Linear detector estimation. Kalker [98] proposes an attack based on the assumption that a public watermark detector is available. A public watermark detector is a device that is freely available and outputs a yes/no answer to indicate the presence of a watermark given a certain image. Kalker's analysis shows that given such a detector device, it is computationally feasible to break watermarking schemes based on linear additive embedding.

Smooth area analysis. Fridrich's own analysis [71] shows that in reasonably large smooth image regions, it is feasible to estimate the embedded watermark given the knowledge about the underlying image transform and assuming a linear, additive embedding method.

In many papers, the authors explicitly call for a security framework that conceals the structure of the transform to improve security. However, few methods are known to construct a key-dependent transform which obeys important properties such as energy-compaction and approximate HVS modeling. Besides Fridrich's approach described in detail in the next section, Ramkumar [195, 197] proposes to scramble selected coefficients with an invertible cyclic all-pass filter (depending on a secret parameter). However, this adds an additional processing step and has therefore higher computational requirements.

4.2 Key-dependent Basis Functions

Watermarking schemes that embed the watermark in certain significant DCT coefficients have been shown to be very robust [39]. The embedding process can also be seen as a modification of the host image's projection onto smooth orthogonal basis functions – discrete cosines in the case of the DCT.

Fridrich [71] demonstrated a hostile attack that exploits the knowledge of Cox's [39] watermarking algorithm to reconstruct the hidden watermark sequence given the marked image. Once the secret watermark sequence is revealed, the watermark can be simply subtracted from the image. The attack is based on two assumptions:

1. The host image contains regions whose unwatermarked pixel values are known or can be easily guessed. Especially in smooth regions of uniform brightness or uniform gradient, the original, unwatermarked pixel values can be accurately estimated.

2. The basis vectors of the transform are publicly known, as it is the case for the common image transforms, such as the DCT or the wavelet decomposition with known filters.

Under these assumptions, it is possible to set up a system of linear equations whose solution allows to determine the embedded watermark. However, in Fridrich's experiment the watermarking schemes was limited to modify only 50 instead of 1000 DCT coefficients. Nevertheless, the proposed attack suggests current watermarking schemes might be vulnerable especially in smooth image regions.

4.2.1 Algorithm Fridrich

Authors This algorithm has been developed by Jiri Fridrich at the Central for Intelligent Systems, SUNY Binghamton, NY, USA and Arnold C. Baldoz and Richard J. Simard at the Air Force Research Laboratory, Rome, NY, USA and is published in [71, 70].

Watermark The mark w_i is a sequence of binary values with length N.

Decomposition No decomposition is used. The image is projected onto the N pseudo-noise patterns P_i. The patterns P_i are generated in a key-dependent pseudo-random way. Then the pattern are smoothed by employing an averaging filter. This step is necessary to achieve robustness and invisibility because most energy will be placed in low-frequency regions. Next, the Gram-Schmidt algorithm is applied to orthogonalize the patterns.

Coefficient selection The entire image is manipulated.

Embedding First, the host image in converted to an intensity matrix, $F(m,n) \in [0,1]$. Then a smoothing filter is applied several times to concentrate the energy in the low-frequency components. Finally, the mark w_i is embedded using the additive embedding formula

$$F'(m,n) = F(m,n) + \alpha \cdot \sum_{i=0}^{N-1} w_i \cdot p_i \cdot P_i,$$

where P_i is a pseudo-random and smoothed pattern and p_i is a projection of the host image F onto the pattern P_i,

$$p_i = P_i \cdot F.$$

Extraction The inverse embedding formula is used to extract the watermark.

Discussion Due to the high computational complexity and storage requirements, the original algorithm was almost impractical except for a very small image size. In a later work [70], the complexity could be considerably reduced.

(a) watermarked image

(b) difference image

Parameter	Description
$\alpha = 0.15$	quantization step size
$N = 64$	watermark length

Figure 4.1: Watermarked image (a) and difference image (b), created with Fridrich's algorithm and the embedding parameters above.

4.3 Parametrization of Wavelet Filters

In this section, we will focus on the possibility to construct secret wavelet filters to improve the security of watermarking applications. Fridrich [71, 70] introduced the concept of key-dependent basis functions in order to protect a watermark from hostile attacks. Hostile or intentional attacks exploit the knowledge of the watermarking algorithm to destroy or remove the watermark [262]. By embedding the watermark information in a secret transform domain, Fridrich's algorithm can better withstand attacks such as those described by Kalker [97] employing a public watermark detector device and other attacks exploiting information about the watermarking algorithm. However, Fridrich's approach suffers from the computational complexity and the storage requirements for generating numerous orthogonal patterns of the size of the host image. Nevertheless, watermarking schemes such as those presented by Wang [254], Zhu [275], Lin [136] or Xia [266] call for a mechanism to protect the location where watermark information has been embedded.

Other security techniques, such as pseudo-random skipping of coefficients, seriously limit the robustness and capacity of the scheme. Therefore, we propose to construct secret wavelet filters by parametrization to decompose the host image. Due to the secret transform domain, the location of the watermark information is protected. Several parametrizations for orthogonal and bi-orthogonal wavelet filters are readily available [191, 277, 201], allowing to choose parameters from a vast key-space. We will show the applicability of this approach and demonstrate its robustness and security.

4.3.1 Zou's parametrization

In order to construct compactly supported orthonormal wavelets, solutions for the dilation equation

$$\phi(t) = \sum_{k \in \mathbb{Z}} c_k \phi(2t - k),$$

with $c_k \in \mathbb{R}$, have to be derived, satisfying two conditions on the coefficients c_k [47]. Schneid [208] describes a parametrization for suitable coefficients c_k based on the work of Zou [277] to facilitate construction of such wavelets. Given N parameter values $-\pi \leq \alpha_i \leq \pi$, $0 \leq i < N$, the recursion

$$\begin{aligned}
c_0^0 &= \tfrac{1}{\sqrt{2}} \text{ and } c_1^0 = \tfrac{1}{\sqrt{2}} \\
c_k^n &= \tfrac{1}{2}(c_{k-2}^{n-1} + c_k^{n-1}) + \\
&\quad \tfrac{1}{2}(c_{k-2}^{n-1} - c_k^{n-1}) \cos \alpha_{n-1} + \\
&\quad \tfrac{1}{2}(c_{2(n+1)-k-1}^{n-1} - c_{2(n+1)-k-3}^{n-1})(-1)^k \sin \alpha_{n-1}
\end{aligned}$$

can be used to determine the filter coefficients c_k^N, $0 \leq k < 2N + 2$. We set $c_k = 0$ for $k < 0$ and $k \geq 2N + 2$.

The parameters α_i are kept secret. These parameters are the key to the wavelet transform domain that can be obtained when decomposing an image with the constructed wavelet filters. The above parametrization of wavelet filter coefficients generates perfect reconstruction filters. Figure 4.2 illustrates what some particular parametric wavelets of filter length 6 look like compared with the Daubechies-6 wavelet.

4.3.2 Pollen's parametrization

Pollen [191] proposed a parametrization for constructing 6-tap orthogonal filters. Two parameters $-\pi \leq \alpha, \beta < \pi$, which are kept secret, are used to control the filter construction. The resulting filters are guaranteed to achieve perfect reconstruction.

$$c_{-2} = ((1 + \cos\alpha + \sin\alpha) * (1 - \cos\beta - \sin\beta) + 2 * \sin\beta * \cos\alpha)/4$$

$$c_{-1} = ((1 - \cos\alpha + \sin\alpha) * (1 + \cos\beta - \sin\beta) - 2 * \sin\beta * \cos\alpha)/4$$

$$c_0 = (1 + \cos(\alpha - \beta) + \sin(\alpha - \beta))/2$$

$$c_1 = (1 + \cos(\alpha - \beta) - \sin(\alpha - \beta))/2$$

$$c_2 = 1 - c_{-2} - c_0$$

$$c_3 = 1 - c_{-1} - c_1$$

4.3.3 Application to Watermarking

We propose to decompose the host image for watermarking purposes using wavelet filters constructed with one of the above parametrizations. The parameter values used for construction and the resulting wavelet filter coefficients are kept secret. Hence, the watermark information can be embedded in a secret

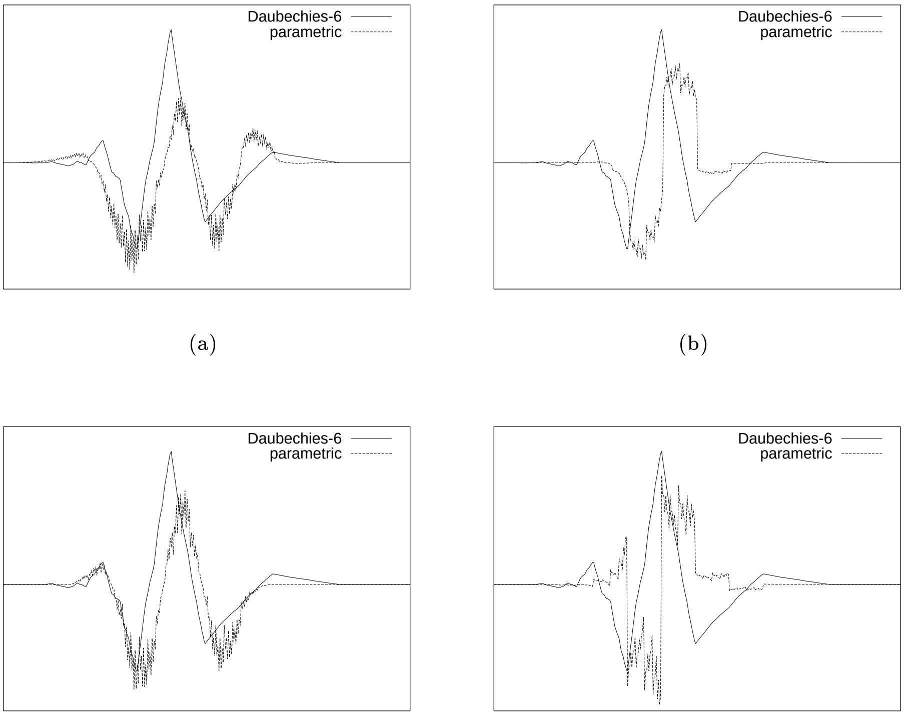

(a) (b)

(c) (d)

Figure 4.2: In each figure, the Daubechies-6 wavelet is compared with a parametric wavelet (Zou's parametrization) of the same filter length, constructed with the following parameters: ($\alpha_0 = -0.6615$, $a_1 = 2.9085$) (a), ($\alpha_0 = -0.0715$, $a_1 = 3.0585$) (b), ($\alpha_0 = -0.4815$, $a_1 = 2.6585$) (c), ($\alpha_0 = 0.1185$, $a_1 = -0.0115$) (d).

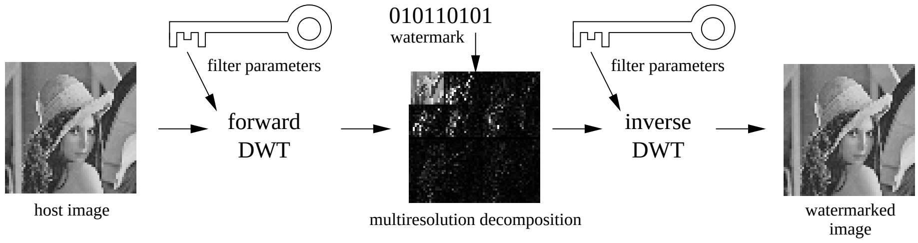

Figure 4.3: Illustration of the watermark embedding process based the concept of key-dependent wavelet filters.

multi-resolution transform domain, making it difficult to mount a hostile attack that seeks to destroy or remove watermark information at specific locations. Our concept is illustrated in figure 4.3.

4.3.4 Decomposition Properties

A problem with randomly-constructed parametric wavelet filters is that the high-pass/low-pass decomposition property is partially lost. Some degree of wavelet smoothness is desirable for most applications. Therefore, we calculate the second-order local variation (difference) of a wavelet sequence

$$V_{\phi}^{(2)} = \sum_{n} \left| c_n^{(J)} - c_{n-1}^{(J)} + c_{n-2}^{(J)} \right|$$

as a simple measure to ensure wavelet smoothness [155]. We can restrict our key-space to parameters such that only wavelets of certain smoothness are produced, e.g. $V_{\phi_\alpha}^{(2)} < V_H^{(2)}$, where $V_H^{(2)}$ is the smoothness measure of the Haar wavelet. Clearly, this is a tradeoff between the security (key-space) and the desireable decomposition properties of the transform.

Hsu [81] states that the choice of the wavelet filter is a critical issue for the quality of the watermarked image and the robustness to compression attacks. However, the filter criteria for watermarking purposes are different compared to image compression applications. Filters that pack most energy of the original image in the lowest resolution approximation image give best compression performance because information in the detail subbands can be easily discarded without severe perceptible image distortion. However, watermarking applications using such filters to embed watermark information in the detail subbands will seriously suffer from compression attacks. Currently, the suitability of different transform domains and wavelet filters are evaluated for watermarking applications with regard to image compression attacks [121].

Employing secret filter parametrization in wavelet-based watermarking algorithms has the following advantages.

Security is improved because unfriendly attacks have to operate in the transform domain used for watermark embedding. Our experiments indicate that the size of the key-space is at least 63000 parameter combinations.

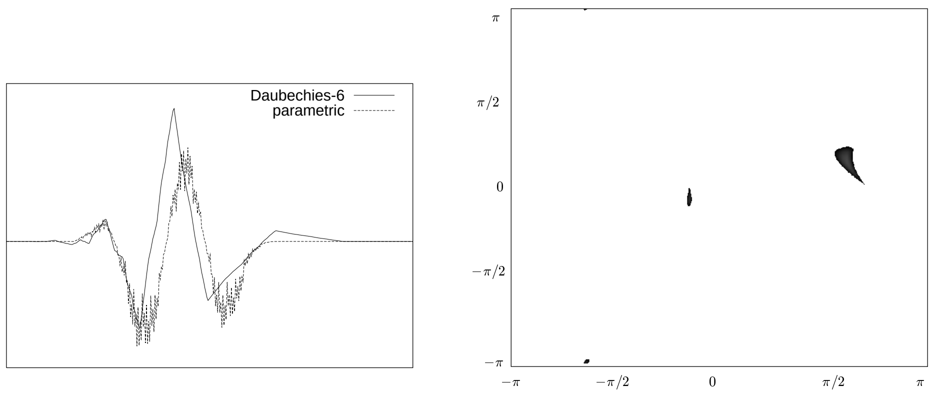

(a) parametric wavelet (b) regions of smoothness

Figure 4.4: The Daubechies-6 and a parametric wavelet ($\alpha_0 = -0.4815$, $a_1 = 2.6585$) (a) and its regions of smoothness (black) with $V_{\phi_\alpha}^{(2)} < V_H^{(2)}$ for $N = 2$ (b).

Image-adaptive. Filter coefficients for watermark embedding can be constructed in an image-adaptive way to maximize robustness against specific compression attacks.

No modification. There is no need to modify proven watermarking schemes (only absolute thresholds have to be adjusted).

Efficient. Furthermore, there is no additional computational cost as the code for the wavelet transform does not have to be changed, only the decomposition filters are generated according to the secret keys.

A wavelet transform based on secret filters can therefore act as a security framework independent of the embedding algorithm.

4.3.5 Results

We conduct all our experiments with the 512×512 gray-scale image 'Lena'. One blind and two non-blind wavelet-based watermarking algorithms (by Kundur, Wang and Kim, described in chapter 3) are used to embed and extract watermark information without perceptible image degradation. The performance of the watermarking schemes is evaluated by calculating the normalized correlation measure.

4.3.5.1 Robustness

First, we demonstrate the robustness against compression attacks that can be achieved when using randomly chosen wavelet filter parameters. We construct 169 different wavelet filters, uniformly separated in the parameter space

($N = 2$; $\alpha_0, \alpha_1 \in \{-3.1, -2.6, \ldots, 2.4, 2.9\}$; $\Delta = 0.5$). Next, we embed a watermark in the host image using one of the available parametric filters for wavelet decomposition; for reference we also test the Daubechies-6 and 9/7-bi-orthogonal filter. The embedding algorithms are discussed in detail in chapter 3. The watermarked images are subjected to JPEG and JPEG2000[1] compression with different quality or bit-rate settings, respectively, resulting in compression ratios from approximately 1 : 4 up to 1 : 80. Figures 4.5 and 4.6 show that all wavelet filters provide adequate robustness, however, the 9/7-bi-orthogonal filter gives best results. We conducted the experiment with all 169 parametric filters but only show the average correlation. The performance of our parametric filters can be improved by restricting the parameter space such that only reasonable smooth wavelets are used. In that case, one can expect results for the parametric filters that are close to the Daubechies-6 filter, compare with figure 4.4.

4.3.5.2 Security

The next experiment examines the security of our filter parametrization approach. For each algorithm, we generate a watermark and embed it using a secret parametric wavelet filter (e.g. $\alpha_0 = 1.7585, \alpha_1 = 1.0585$). Then we try to extract that watermark but randomly 'guess' the transform parameters within the key-space. Figures 4.7 and 4.8 suggest that the watermark can be retrieved correctly only with matching wavelet filters. The "Lena" image was watermarked with the algorithms by Kundur, Kim and Wang using a particular key to construct the wavelet filters. First, we generate 3969 distinct keys together with the corresponding wavelet filters and compute the correlation between the embedded and the extracted watermark, employing each prepared decomposition filters. The normalized correlation result is depicted in figure 4.7 for each watermarking scheme. In the second case, we tested 63504 uniformly distributed parameters ($N = 2; \alpha_0, \alpha_1 \in \{-3.14, -3.11, \ldots, 3.11, 3.13\}; \Delta = 0.025\}$), see figure 4.8. In addition, the normalized correlation result for 63×63 parameters, uniformly distributed in the key space, is shown as a "map" where the bright regions indicate high correlation; see figure 4.9.

We repeat the experiment but restrict the key-space to parameters that produce smooth wavelets according to our measure, $V_{\phi_\alpha}^2 < V_H^2$. The embedded watermark can only be retrieved with matching parametric filters, see figure 4.10.

One limitation of our security approach can be seen with the watermarking algorithm proposed by Kundur. The watermark can be extracted even when the extraction key is only "close" to the embedding key. This result is due to the quantization embedding strategy of that algorithm which rejects small amounts of distortion in the image data and the limited length of the binary watermark sequence. However, further results[2] suggest that the concept of parametric wavelet filters is applicable for both, quantization and additive embedding methods.

[1] Using JasPer (based on the JPEG2000 working draft), see `http://spmg.ece.ubc.ca/people/mdadams/jasper/index.html`.

[2] Available at `http://www.cosy.sbg.ac.at/~pmeerw/Watermarking`.

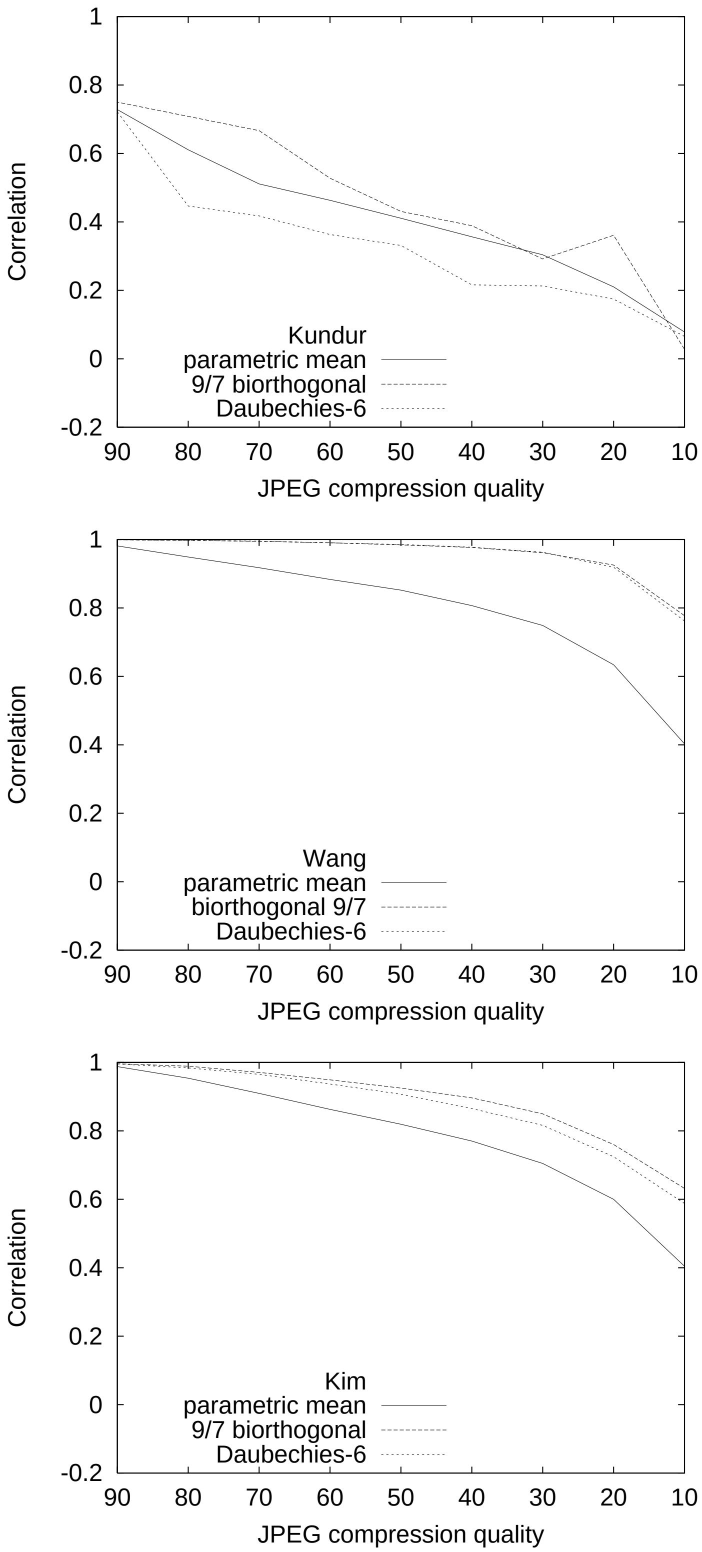

Figure 4.5: The mark was embedded using parametric, Daubechies-6 and 9/7-bi-orthogonal filters with the watermarking algorithms by Kundur, Wang and Kim. The correlation results of the extracted watermark after JPEG attack for the different algorithms is shown in first through third row.

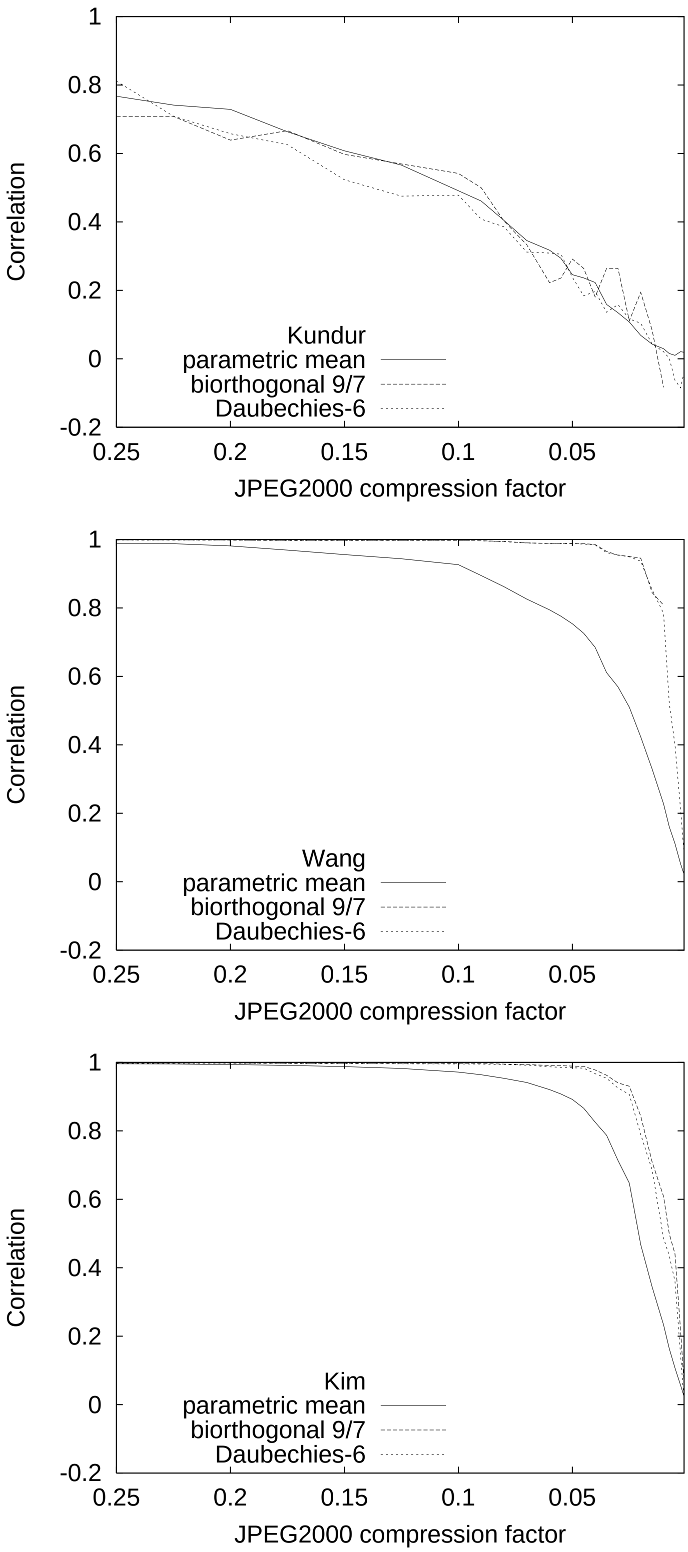

Figure 4.6: The mark was embedded using parametric, Daubechies-6 and 9/7-bi-orthogonal filters with the watermarking algorithms by Kundur, Wang and Kim. The correlation results of the extracted watermark after JPEG2000 attack for the different algorithms is shown in first through third row.

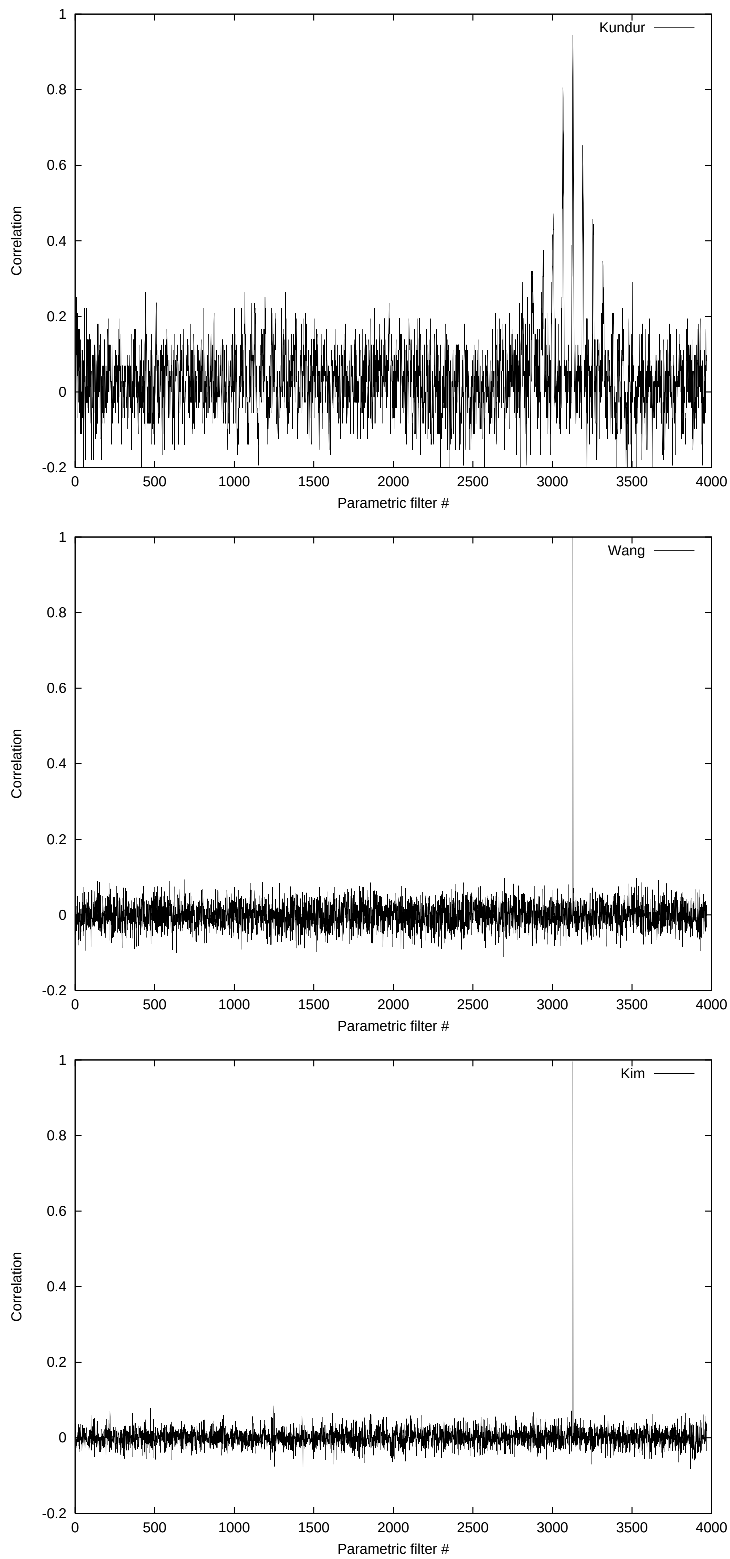

Figure 4.7: The security of 4000 parametric key-dependent wavelet filters. The embedded watermark can only be detected with matching keys and wavelet filters.

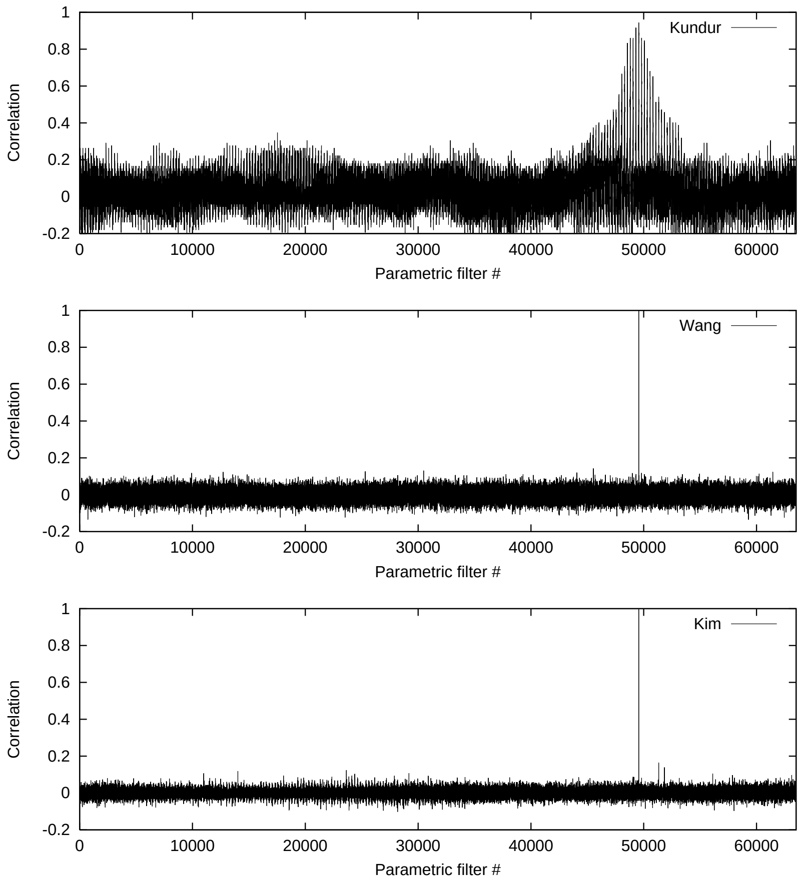

Figure 4.8: The security of 65000 parametric key-dependent wavelet filters. The embedded watermark can only be detected with matching keys and wavelet filters.

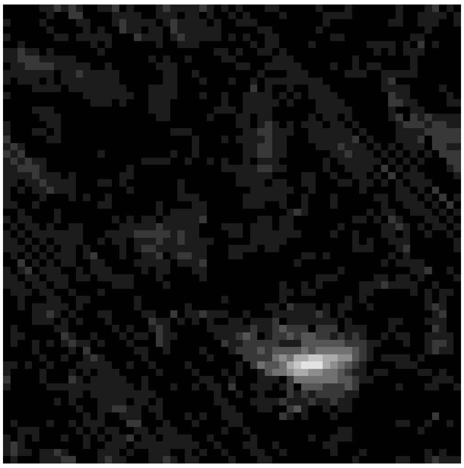

(a) Kundur

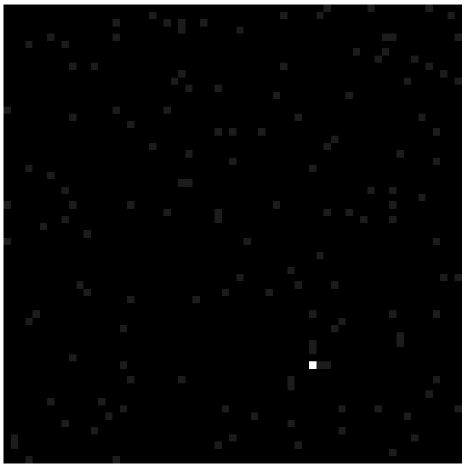

(b) Wang

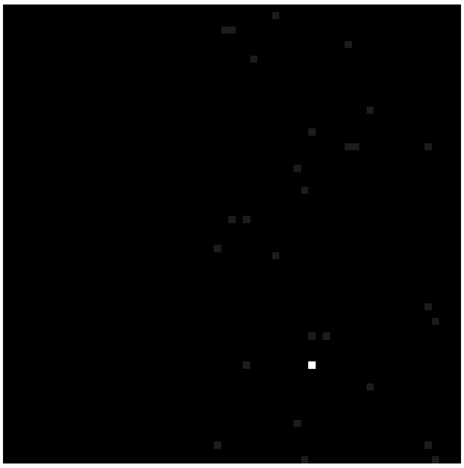

(c) Kim

Figure 4.9: Correlation map for the parameters within the key-space. Computing the normalized correlation between the embedded and the extracted watermark for 63×63 parameters using Zou's wavelet filter parametrization to contruct 6-tap filters. The watermarking algorithm by Kundur, Wang and Kim were used for the experiments.

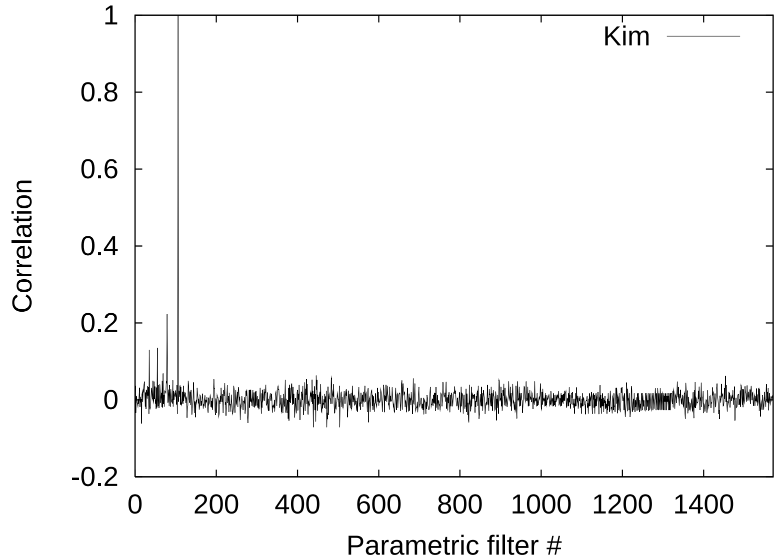

Figure 4.10: The security of smooth parametric key-dependent wavelet filters (Zou's parametrization, 6-tap filter) according to our smoothness measure. The embedded watermark (Kim's algorithm) can only be detected with matching keys and wavelet filters.

4.4 JPEG2000 integrated watermarking

An early attempt to integrate wavelet-based image coding and watermarking has been made by Wang [253, 254] and Su [224]. While the first approach was based on the "Multi-Threshold Wavelet Codec" (MTWC) [251, 250], the later proposal builds on "Embedded Block Coding with Optimized Truncation" (EBCOT) [232] which is also the basis for the upcoming JPEG2000 image compression standard. Both watermarking algorithms add a pseudo-random Gaussian noise sequence to the significant coefficients of selected detail subbands.

In this section, we present a blind watermarking technique integrated in the JPEG2000 coding pipeline. The watermark embedding and recovery process is performed on-the-fly during image compression and decompression. The computational cost to derive the transform domain a second time for watermarking purposes can therefore be saved.

Our design builds on the results of the previously proposed wavelet-domain watermarking algorithms mentioned above. However, in order to fit the JPEG2000 coding process, our watermarking system has to obey the independent processing of the code-blocks. Algorithms which depend on the inter-subband [111] or the hierarchical multi-resolution [86] relationship can not be used directly in JPEG2000 coding. Due to the limited number of coefficients in a JPEG2000 code-block, correlation-based methods [4, 253] fail to reliably detect watermark information in a single independent block. Obviously, watermarking methods that require access to the original image or reference data for watermark extraction are not suited as well – this precludes all the non-blind schemes

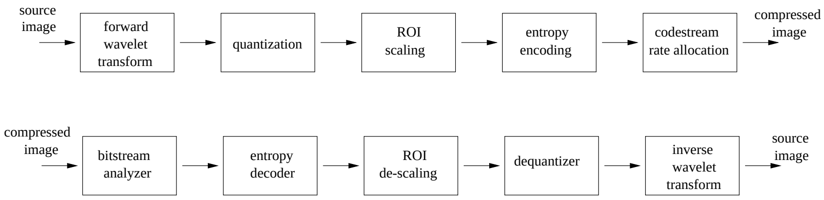

Figure 4.11: The JPEG2000 coding pipeline.

[34, 266, 104].

Two application scenarios are demonstrated: copyright protection and image authentication ("tamper detection").

The main design goals behind EBCOT and JPEG2000 are versatility and flexibility which are achieved to a large extent by the independent processing and coding of image blocks [27]. The default for JPEG2000 is to perform a five-level wavelet decomposition with 7/9-biorthogonal filters and then segment the transformed image into non-overlapping code-blocks of no more than 4096 coefficients which are passed down the coding pipeline.

4.4.1 Watermark Embedding

The watermark embedding stage is invoked after quantization and region-of-interest (ROI) scaling and prior to entropy coding (see figure 4.12). At that point, each code-block transports signed integer coefficients that have been normalized: the most significant bit (MSB) carries the sign bit and the remaining bits represent the absolute magnitude of the coefficient (the actual number depends on the implementation, we assume 32 bits for this work).

We have to distinguish between code-blocks belonging to either the approximation image (LL subband) or the detail subbands (LH_j, HL_j, HH_j subbands, where $j = 1 \ldots J$ is the decomposition level). The finest resolution subbands can not be used to encode information reliably.

In the first case, i.e. code-blocks belonging to the approximation image, we apply an embedding technique similar to Xie's [268] approach. We slide a non-overlapping $w \times 1$ running window over the entire code-block. At each window position, one bit of watermark information is encoded using the quantization embedding technique described in section 3.4. The size of the embedding window determines the coding rate. Given a gray-scale image of size 512×512, the watermark information that can be embedded in the approximation image is $\frac{512 \cdot 512}{2^{2 \cdot J}} \cdot \frac{1}{w}$ bits. Typical values of w range from 2 to 8. We found that the quantization step size Δ has to be chosen around 2 500 000 for imperceptible yet robust embedding.

For the code-blocks being part of one of the detail subbands, we have to use a larger embedding window because the energy is much lower. At least 256 coefficients are quantized to encode one watermark bit. Thus, if the size of the code-block allows, we can split it into several sub-blocks to increase the embedding capacity. The large magnitude coefficients represent edge and texture

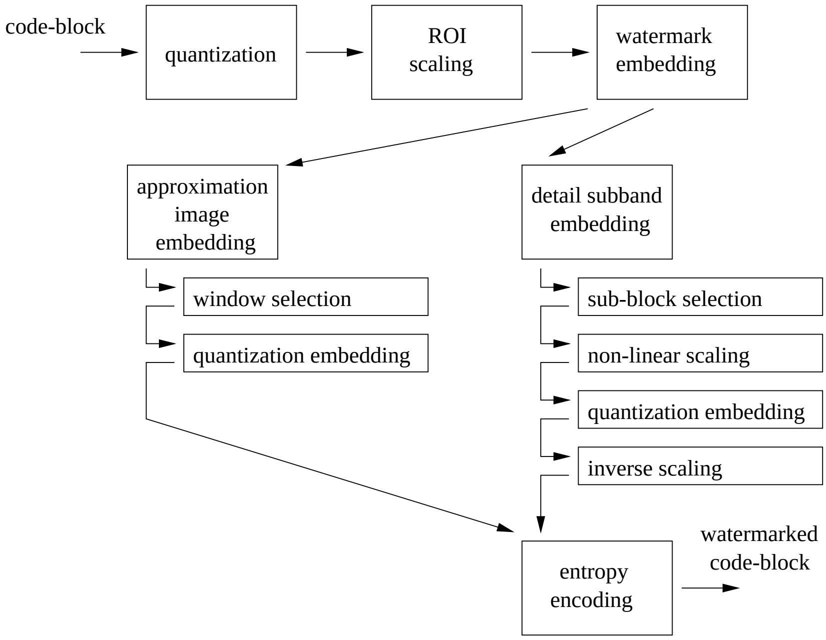

Figure 4.12: The watermark embedding process in the JPEG2000 coding pipeline.

decomposition level j	scaling factor β
2	6.5
3	6
4	5.5
5	5

Table 4.1: Decomposition-level dependent parameter for the non-linear scaling function.

information. The human visual system (HVS) is less sensitive to changes in these regions, therefore we want to exploit this characteristic to maximize the watermark strength. To keep the implementation simple, a non-linear scaling function $f(x) = sign(x) \cdot \left|x^{\beta}\right|$, $\beta > 1$ is applied to all code-block coefficients. The scaling parameter β is chosen in a level-adaptive way according to table 4.1. We obtain a more uniform coefficient representation since the high peaks in the coefficient distribution are reduced. This way, we can use simple uniform scalar quantization (as before) and still put more watermark energy in the image regions the HVS is less sensitive to. After quantization, the inverse scaling function f^{-1} is applied to derive the watermarked code-block.

image	window size	sub-block size	capacity	PSNR
Lena	4	64	85	32.05
Fishing Boat	2	32	194	31.45
Goldhill	2	16	383	32.09

Table 4.2: Embedding parameters and the corresponding bit capacity for three 512×512 gray-scale images, together with the resulting PSNR.

4.4.2 Results

We conducted our experiments with the JJ2000 (version 3.2.2) implementation[3] of the JPEG2000 verification model (VM). The modularized architecture of the JJ2000 software allowed to easily integrate our watermarking module. If not noted otherwise, we use the default coding parameters for our experiments. This means that the 7/9-biorthogonal wavelet filters are employed to decompose the host image into a five level multi-resolution representation.

Our watermarking method was tested in two application scenarios: copyright protection and image authentication (or tamper detection).

4.4.2.1 Copyright protection

To demonstrate the robustness and capacity of our watermarking method, a watermark with 85, 194, and 383 bits was embedded in the 512×512 gray-scale images "Lena", "Fishing Boat" and "Goldhill", respectively. Figure 4.13 shows the watermarked images "Lena" and "Goldhill" on the left and their difference images on the right side. The different watermark capacities were achieved by choosing the embedding parameters from table 4.2. The resulting PSNR is also given. The effect of our simple scaling function is clearly visible in the difference images: the edges contain more watermark energy than smooth regions.

For copyright protection, we embed a binary message that identifies the owner of the image. The dither vectors are kept secret to protect the watermark. The normalized correlation result of the recovered versus the embedded message is depicted in figure 4.15. The watermarked images were subjected to JPEG and JPEG2000 compression with varying compression parameters (top row). To simulate image processing attacks, the images were blurred and sharpened using the ImageMagick[4] convert program (bottom row). The results indicate our watermark survives the attacks, but additional error-corrective coding is required to achieve perfect recovery of the embedded information.

4.4.2.2 Tamper detection

The tamper detection application requires a fragile watermark that breaks in order to indicate the areas that have been manipulated. At the same time, however, the watermark should be robust against unintentional distortion, e.g. caused by lossy image compression. Figure 4.14 (a) shows the "Fishing Boat"

[3] The JJ2000 source is available for download at `http://jj2000.epfl.ch`.

[4] The ImageMagick programs are at `http://www.simplesystems.org/ImageMagick`.

image, watermarked with a sequence of all-zero bits. Next, we manipulated three regions using the GIMP[5] and JPEG compressed the image with default quality; see figure 4.14 (b) and the difference image, highlighting the changes (c). The detection results of our watermarking schemes are depicted in figure 4.14 (d). The malicious tampering has been detected and localized while the distortion due to JPEG compression did not raise a false alarm.

One coefficient in the approximation image of the wavelet domain corresponds to a block of pixels in the spatial domain. In order to achieve good spatial resolution for our tamper detection example, we had to limit the wavelet transform to three decomposition steps. Therefore, we can authenticate pixel blocks of size 8×8 individually. Since the watermark consists of sequence of zero bits, we could use sliding window detection [60] in horizontal, diagonal and vertical direction in the approximation image. The tamper detection results from the three directions were accumulated and contribute to the brightness of the tamper detection image of figure 4.14 (d).

4.5 Conclusions

We have introduced the concept of wavelet filter parametrization to improve the security of watermarking applications. Our approach is easy to integrate in existing watermarking schemes. The experiments indicate that the level of security provided is adequate for many applications. Because our proposed security framework does not require any computational overhead, it is especially suited for video watermarking or other real-time applications. Further work will investigate the parametrization of bi-orthogonal wavelet filters.

We demonstrated that watermarking can be integrated in the JPEG2000 coding process and discussed some of the limitations. A novel embedding algorithm based on QIM and suitable for watermarking independent JPEG2000 code-blocks was proposed which allows blind watermark recovery during image decompression. We investigated a copyright protection and an image authentication application and provided robustness as well as capacity results. Future work will try to improve the performance of the embedding method and consider ROI coding and color images.

[5] The GNU Image Manipulation Program is available at `http://www.gimp.org`.

(a) watermarked image (Lena)

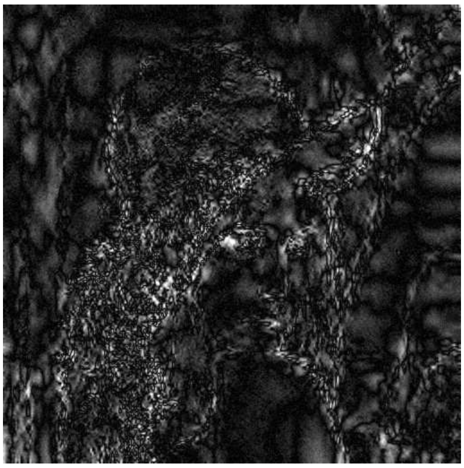

(b) difference image

(c) watermarked image (Goldhill)

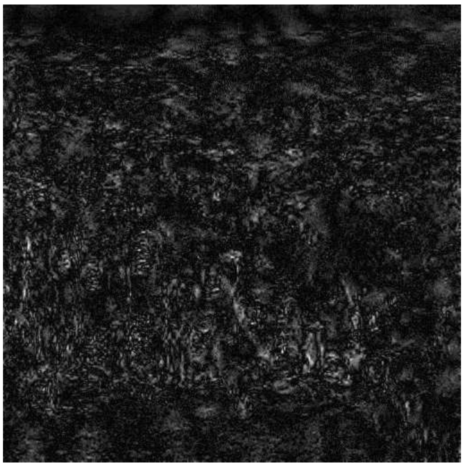

(d) difference image

Figure 4.13: The watermarked images “Lena” (a) and “Goldhill” (c) and their difference images (b) and (d).

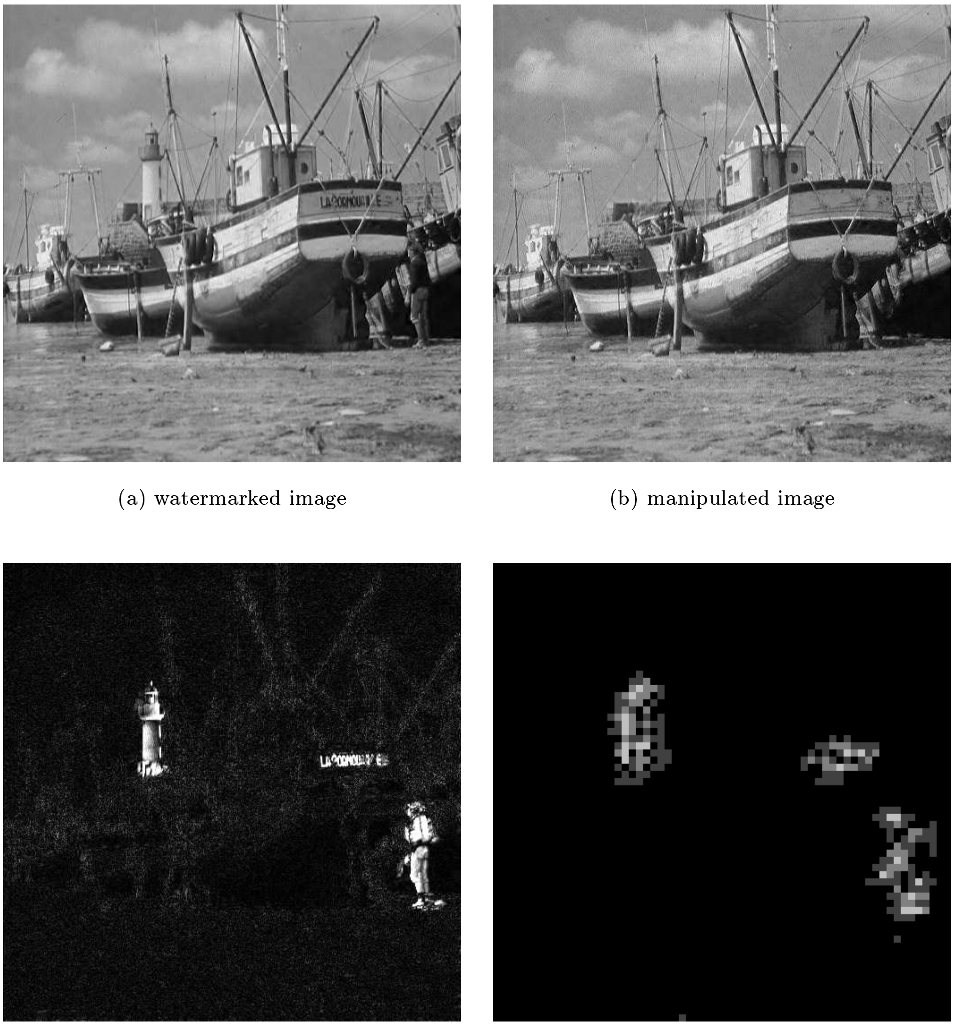

(a) watermarked image

(b) manipulated image

(c) difference image

(d) tamper detection

Figure 4.14: The watermarked "Fishing Boat" image (a) and the tampered version (b). The manipulations are highlighted after default JPEG compression in the difference image (c). The manipulated regions detected by the algorithm (d).

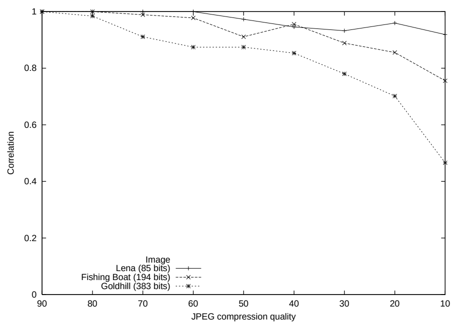

(a) JPEG

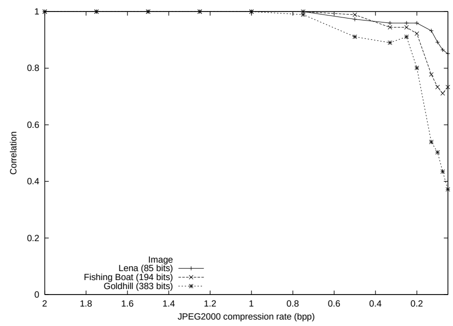

(b) JPEG2000

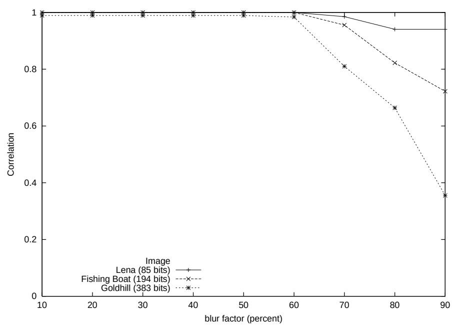

(c) blur

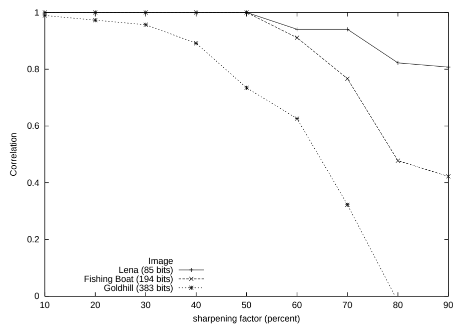

(d) sharpen

Figure 4.15: The robustness against JPEG compression (a) and JPEG2000 compression (b), blurring (c) and sharpening (d). Measured using normalized correlation between the recovered and embedded message.

Chapter 5

Attacks

So far, security issues have been studied in a large number of contributions. One of the first attacks, reported by Craver [43], aims at the ownership claim established with the help of non-blind linear watermarking algorithms (such as Cox [39]). Here, the invertibility property of the embedding formula (see section 3.3) is exploited to create a secondary "fake" watermark which can not be distinguished from the legitimate mark. Thus, two ownership claims, based on one legitimate and one illegitimate watermarked, can not be resolved.

Other attacks have been described by Stone [216] (pseudo-noise attacks), Hartung [77] (spread-spectrum attacks), Kilian [102] (collusion attacks), Kalker and Linnartz [93, 94, 98, 97, 140] (linear detector analysis), Su [221, 217, 218] (distortion bound attack analysis), Eggers [59] (quantization attacks) , Memon [158] (fragile watermark attack), Kutter [128] (copy attack), Dugelay [56] (geometric counter-attack), Voloshynovskiy [246, 243] (noise-removal attack), Fridrich [71, 70] (attack against known image regions), Wu [264] (block interpolation).

In the following sections, we want to provide an overview of the different kinds of attacks (section 5.1) against watermarked image data and the counter-measures (section 5.2) that can be taken to make the embedded mark more secure and robust.

StirMark is a very successful and publicly available software[1], that can be used to "benchmark" watermarking systems [127, 178].

In chapter 6, we test some of the proposed watermarking algorithms and present robustness results.

5.1 Attack Classification

First of all, we have to distinguish two "reasons" or "purposes" for an attack against a watermark image:

[1] StirMark is available at `http://www.cl.cam.ac.uk/~fapp2`.

- hostile or malicious attacks, which are an attempt to weaken, remove or alter the watermark and
- coincidental attacks, which can occur during common image processing and are not aimed at tampering with the watermark.

Lossy image compression is considered the most common form of attack a watermarking scheme has to withstand. The harsh term “attack” can be easily justified: an efficient image compression has to suppress or discard perceptually irrelevant information – the invisible watermark. This duality has already been noted in section 2.7.2.

The following types of attacks can be invoked to penetrate a watermarking system.

Simple attacks. These attacks does not attempt to isolate the watermark. The goal is to add distortion to the host image in order to render the watermark un-detectable or un-readable. The attack is successful if the watermark can not be detected anymore but the image is still intelligible and can be used for a particular determined purpose. Many such attack operations have been proposed:

- lossy image compression,
- addition of (Gaussian) noise,
- median filtering and blurring,
- re-sampling and re-scaling.

Detection-disabling attacks. These attack attempt to break the correlation detection between the extracted and the original watermark sequence. This can be accomplished by “shuffling” the pixels. The values of corresponding pixels in the attacked and the original image are essentially the same, however, the location has changed. We can distinguish

- geometric attacks, where the image is subjected to translation, rotation, scaling and/or cropping,
- jitter attacks or synchronization attacks that prevent the watermark locations from being found (e.g. by removal and insertion of pixel rows or columns, and
- StirMark attacks. StirMark is a program that can be used to apply may different types of attack. One specific attack introduces non-linear distortion via random “bending” into the image. The effect can be best explained visually (figure 5.1).

Ambiguity attacks or deadlock attacks discredit the authority of the watermark, e.g. by embedding at least on additional watermark [43] or the possibility to copy the watermark from one image to another [128] without control of the legitimate owner of the watermark.

Removal attacks attempt to separate and remove the watermark. Example for this attack technique are

- the collusion attack first discussed by Cox [39],
- denoising [243] and
- non-linear filtering [130].

If the watermarking system or protocol makes not only the watermarked image but also additional devices publicly available, the presence of such devices can be exploited.

- Exploiting the presence of a watermark detector. For example, the DVD copy control mechanism depends on a watermark detection device in every consumer player [161, 95, 14]. Kalker [93, 98] and Linnartz [140] found a "sensitivity" attack that can be used to exploit these simple, low-cost devices. Their approach is to create a test image near the detection boundary and then successively change single pixels until the detector response indicates that a particular pixel value has significant influence on the watermark. This way, a set of influential pixels (or sensitive pixels) can be determined that has the largest influence on the detector while introducing low disturbance into the image when manipulated. This process has linear complexity i.e. $O(N)$.
- Exploiting the presence of a watermark inserter. With the presence of a watermark inserter, the difference image between the watermarked and the original image can be easily computed and analyzed. A public watermark inserter is e.g. provided by the DVD system for copy generation management.

Other attacks can exploit weaknesses of the watermarking scheme. In section 4.1 we have already discussed some of these security issues.

5.2 Counter-Attacks

In order to strengthen watermarking schemes, the following precautions against attacks can be taken [77, 50, 102]. Some of these counter-measures such as image registration can be implemented independently of the actual watermarking method.

Image registration. The received image data has to be "mapped" to the original host image in order to determine the locations where the watermark has to be extracted. Imperfect image registration can result from cropping and other geometric attacks or – especially in video watermarking applications – from synchronization problems. Image registration is a minor problem if the original image is available to the watermark receiver.
On the other hand, without reference to the original image, the registration process poses a serious problem [234, 235]. In this case, one can try to estimate the transformations the image has undergone and reverse their effect. Reference points [125] and reference watermarks (see below) have been proposed to aid in this difficult task.
After all, the registration problem can be seen as a separate stage prior to the watermark extraction stage.

Re-correlation. A counter-measure against StirMark ("bending") and other pseudo-random geometric attacks is to split the image into small blocks and try the estimate the local transform [56]; e.g. try all possible combinations of shift, rotation, zoom to maximize the watermark correlation. Needless to say that these techniques are computationally very expensive. The computational complexity can be reduced, however, by embedding a symmetric watermark such as the circular watermark proposed by Solachidis [214] or Licks [135]. Alternatively, a transform domain invariant to certain geometric attacks can be used, e.g. the Fourier-Mellin [159]. The wavelet domain can provide scale-invariance and shift-invariance [143] to some extent.

Information rate. The number of pixels or coefficients which encode one bit of watermark information should not be too low. Clearly, there is a capacity bound on the watermark channel given a particular amount and type of distortion. A considerable number of contributions has analyzed the achievable capacity of watermarking systems [211, 8, 154, 188, 196].
Of course, redundancy and error-corrective coding is also an issue to build reliable watermarking systems.

Strong cryptographical components. The security of many watermarking algorithms depends on pseudo-random number generators; e.g. to produce a Gaussian sequences of real-numbers for spread-spectrum watermarking or to skip coefficients in a pseudo-random manner. Secure image hash functions are employed in many image authentication schemes. Most watermarking algorithms embedding a logo-type watermark de-correlate the logo image before fusing the images. Thus, a secure permutation or mixing system [207, 248] is needed.

Power-adaptive watermarking. The embedded watermark signal should be as strong as possible to survive energy-bound attacks [219, 113, 59]. Hence, a watermarking scheme has to be image-adaptive to place most of the watermark signal's energy in the high energy components of the host image. The wavelet transform offers very good implicit modeling of the HVS and therefore easily allows robust watermarking while achieving image transparency at the same time. However, exploiting explicit perceptual masking, the performance can be further improved.

Attack characterization. Kundur [116, 119, 111] proposes to embed a known reference watermark to estimate the type of attack the image has undergone. In the watermark extraction and correlation stage, the attack estimate is used to weight the recovered watermark. A similar technique is also proposed by Kutter [125] and other authors.

Complementary modulation. To better withstand a wide range of attacks, Lu [150, 145] uses complementary modulation and evaluates both, the positively and negatively modulated watermark in the watermark extraction stage.

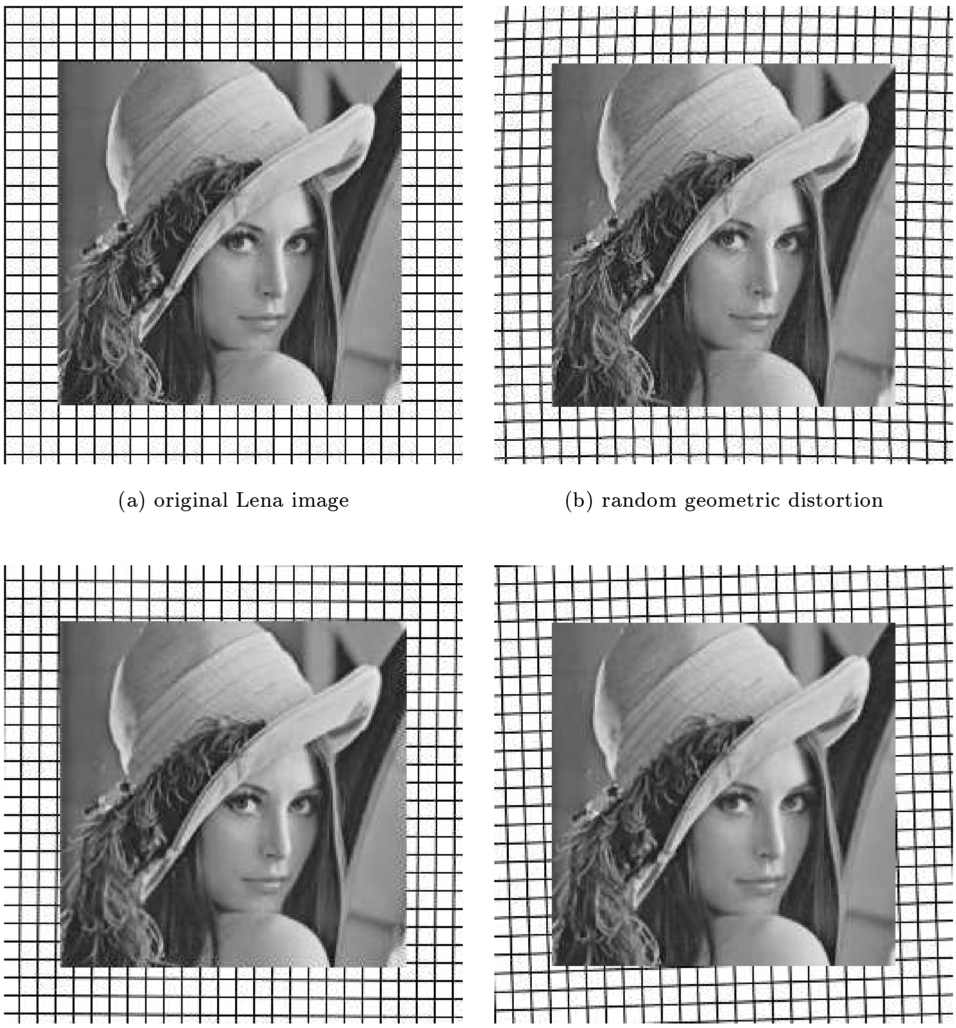

(a) original Lena image

(b) random geometric distortion

(c) rotation by -0.5 degrees

(d) rotation by 2 degrees and scaling

Figure 5.1: Watermark attacks using StirMark. The original Lena image (a). Random bending and JPEG compression (b). Rotation by -0.5 degrees (c). Rotation by 2 degrees and scaled back to the original size (d).

Chapter 6

Results

In this chapter, the measurements performed with the watermarking algorithms that we have implemented are presented. For each algorithm, the embedding parameters were chosen following the description in the papers such that the robustness is maximized. At the same time, the watermarks are still invisible and can not be detected by a human observer familiar to the test images.

The following results are provided:

- the watermark signal strength in PSNR (dB), see section 6.1,
- the watermark capacity in bits, refer to section 6.2,
- the experimental detection threshold (determined with 100 and 1000 random keys), illustrated in section 6.3,
- the normalized correlation between the the embedded and the extracted watermark after the following attacks
  - image compression with JPEG, SPIHT and JPEG2000, shown in section 6.4,
  - image processing (median filtering, smoothing), see section 6.5,
  - geometrical transformation (cropping, down-scaling), see section 6.6.

Furthermore, we illustrate that image registration is a useful counter-measure against StirMark attacks; see section 6.7.

If not noted otherwise, all tests were performed on the 512×512 "Lena" gray-scale image with 8 bits per pixel.

More results are available on my home-page[1].

[1]See `http://www.cosy.sbg.ac.at/~pmeerw/Watermarking`.

6.1 Watermark signal strength

Algorithm	Lena	Baboon	Goldhill	Fishing Boat	Cameraman
Bruyndonckx	42.92	36.79	40.60	42.82	41.66
Cox	38.60	32.05	38.15	34.90	36.52
Corvi	38.23	40.58	48.23	38.85	44.31
Dugad	40.63	27.87	40.74	36.17	39.10
Fridrich	29.51	26.06	25.36	28.65	27.62
Inoue	43.24	39.44	41.00	41.26	39.04
Inoue (insign.)	45.65	46.45	45.15	45.62	41.67
Kim	37.59	31.88	36.15	34.80	35.08
Koch	47.92	45.53	46.89	47.12	41.77
Kundur	48.82	42.82	48.60	47.57	46.64
Wang	33.74	35.30	34.98	34.89	35.78
Wang (blind)	41.45	45.00	41.37	42.78	42.09
Xia	38.52	29.78	33.61	34.15	36.51
Xie	40.12	41.33	43.37	44.42	39.16
Zhu	33.50	35.59	33.91	32.48	36.27

Figure 6.1: The watermark signal strength in PSNR (dB) of selected algorithms and several host images.

6.2 Watermark capacity

Algorithm	Lena	Baboon	Goldhill	Fishing Boat	Cameraman
Bruyndonckx	896	869	890	888	180
Cox	-	-	-	-	-
Corvi	-	-	-	-	-
Dugad	8	9	8	9	9
Fridrich	200	200	200	200	200
Inoue	904	904	904	904	416
Inoue (insign.)	178	184	184	184	110
Kim	-	-	-	-	-
Koch	178	184	174	175	162
Kundur	603	622	618	613	589
Wang	-	-	-	-	-
Wang (blind)	-	-	-	-	-
Xia	-	-	-	-	-
Xie	315	320	290	295	70
Zhu	-	-	-	-	-

Figure 6.2: The watermark capacity in bits of selected algorithms and several host images.

6.3 Detection threshold

6.3.1 Spatial domain algorithms, 100 keys

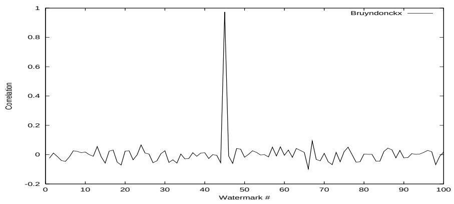

(a) Bruyndonckx

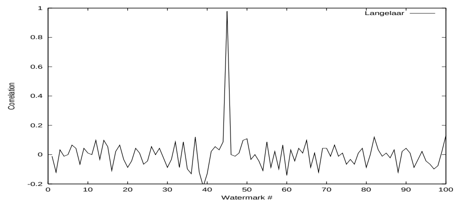

(b) Langelaar

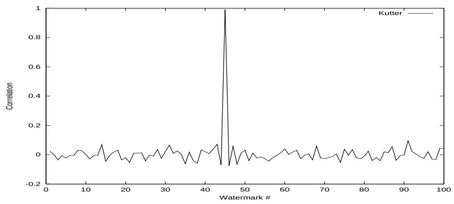

(c) Kutter

Figure 6.3: Correlation of 100 random keys; spatial domain algorithms.

6.3.2 DCT domain algorithms, 100 keys

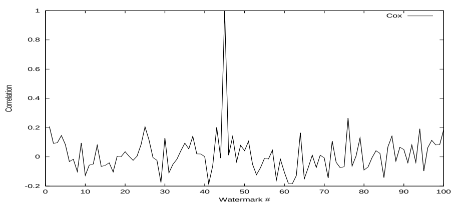

(a) Cox

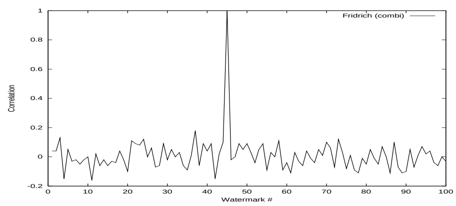

(b) Fridrich

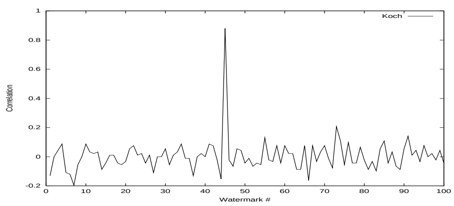

(c) Koch

Figure 6.4: Correlation of 100 random keys; DCT domain algorithms.

6.3.3 Wavelet domain blind algorithms, 100 keys

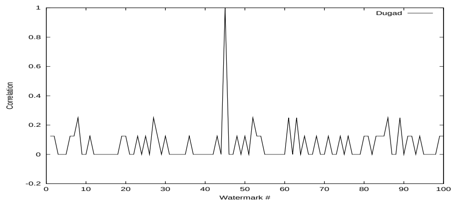

(a) Dugad

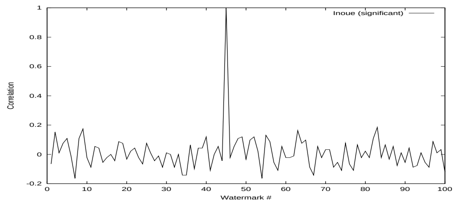

(b) Inoue (insignificant)

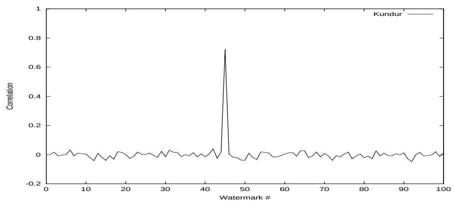

(c) Kundur

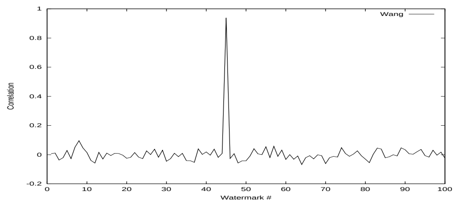

(d) Wang (blind)

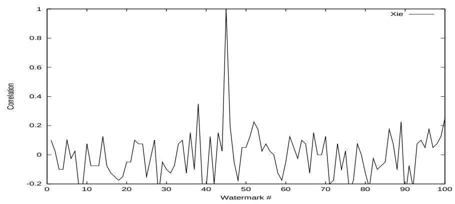

(e) Xie

Figure 6.5: Correlation of 100 random keys; blind wavelet domain algorithms.

6.3.4 Wavelet domain non-blind algorithms, 100 keys

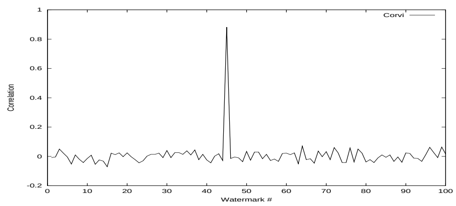

(a) Corvi

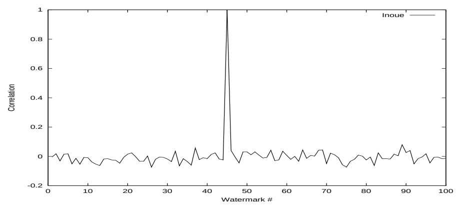

(b) Inoue (significant)

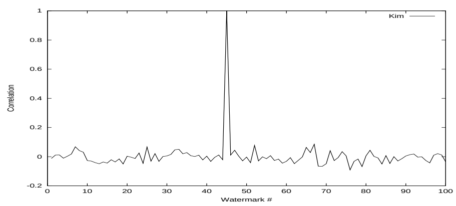

(c) Kim

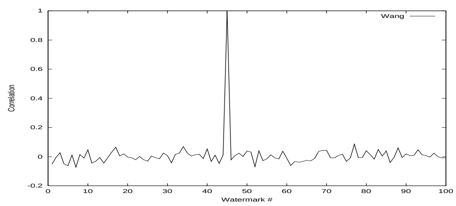

(d) Wang

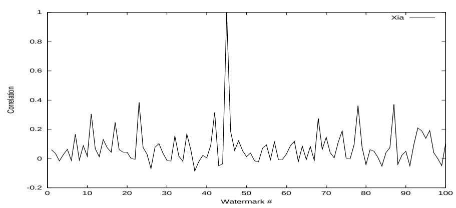

(e) Xia

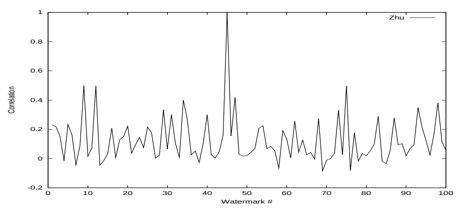

(f) Zhu

Figure 6.6: Correlation of 100 random keys; non-blind wavelet domain algorithms.

6.3.5 Spatial domain algorithms, 1000 keys

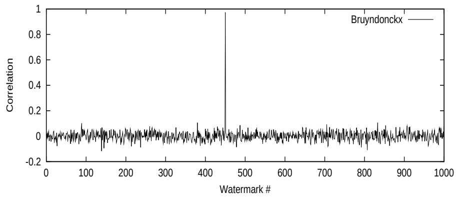

(a) Bruyndonckx

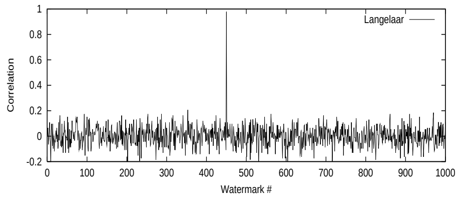

(b) Dugad

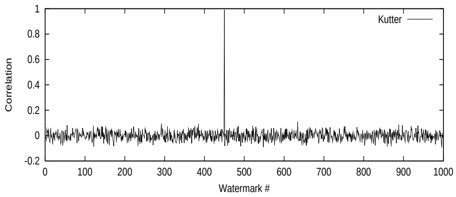

(c) Kutter

Figure 6.7: Correlation of 1000 random keys; spatial domain algorithms.

6.3.6 DCT domain algorithms, 1000 keys

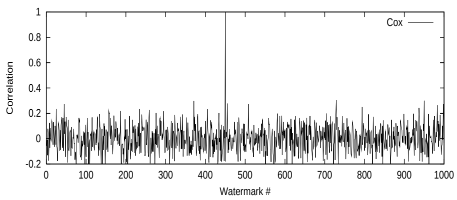

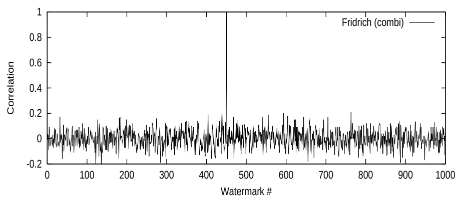

(a) Fridrich

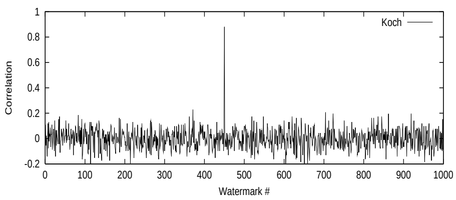

(b) Koch

Figure 6.8: Correlation of 1000 random keys; DCT domain algorithms.

6.3.7 Wavelet domain blind algorithms, 1000 keys

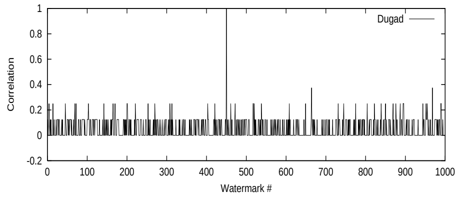

(a) Dugad

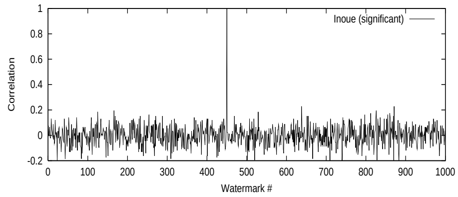

(b) Inoeu (insignificant)

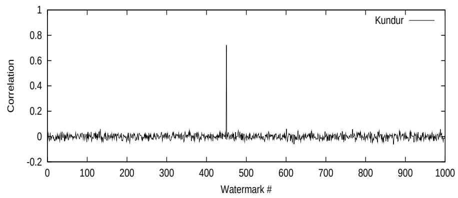

(c) Kundur

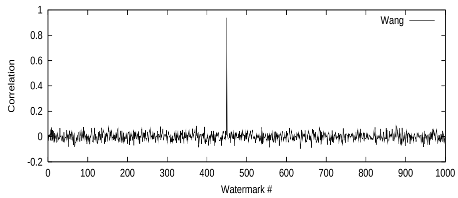

(d) Wang (blind)

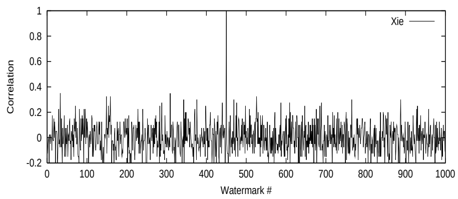

(e) Xie

Figure 6.9: Correlation of 1000 random keys; blind wavelet domain algorithms.

6.3.8 Wavelet domain non-blind algorithms, 1000 keys

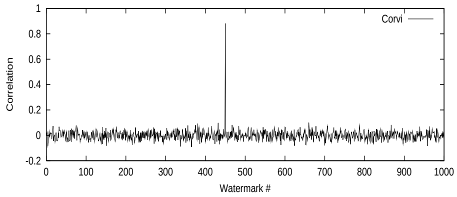

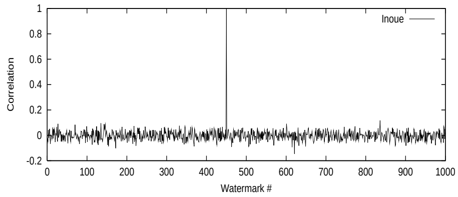

(a) Corvi

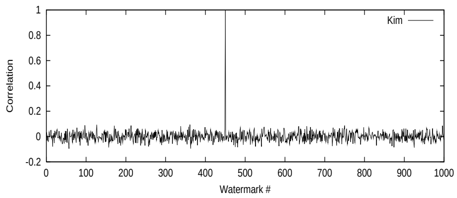

(b) Kim

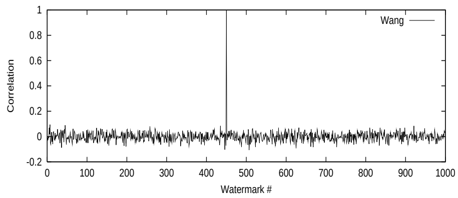

(c) Wang

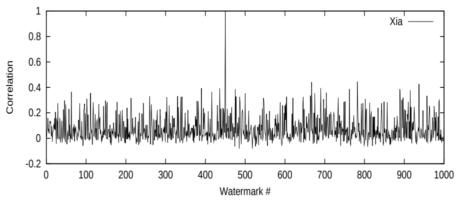

(d) Xia

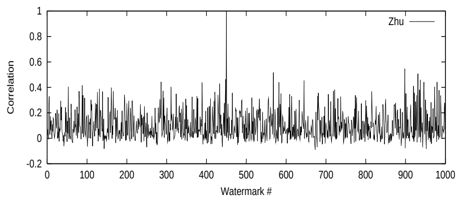

(e) Zhu

Figure 6.10: Correlation of 1000 random keys; non-blind wavelet domain algorithms.

6.4 Image Compression

We test lossy image compression using the JPEG, SPIHT and JPEG2000 coder. In order to compare the different types of distortion caused by the lossy compression schemes and factor out the influence of quality factor versus coding rate settings, we plot the normalized correlation between the embedded and the recovered watermark on a PSNR scale. First, the PSNR between the compressed and the original image was measured for different compression rates with all three coding schemes. Next, the watermarked image was subjected to lossy compression, using the parameters determined before, and the correlation results is recorded together with the associated PSNR. The PSNR range from 43 to 30 dB corresponds to the quality factors of 95 down to 10 for JPEG compression and bit rates of 1.75 to 0.1 per pixel for SPIHT and JPEG2000 compression.

6.4.1 Capacity gap

While lossy image compression systems aim to discard redundant and perceptual insignificant information in the coding process, watermarking schemes try to add invisible information to the image. An optimal image coder would therefore simply remove any embedded watermark information. This duality has been pointed out by a number of authors. However, even state-of-the-art image coding systems such as JPEG2000 [1] do not achieve optimal coding performance and therefore there is a "distortion gap" that can be exploited for watermarking; see figure 6.11. We watermark the "Lena" image with Xie's [268] and Wang's [252] embedding algorithm and plot the rate/distortion performance for both, the original and watermarked image; see (a). The amount of watermark information, measured as normalized correlation, that "survives" the attack is shown as well and demonstrates that the watermark can be recovered until the "capacity gap" closes; compare with figure 6.11 (b).

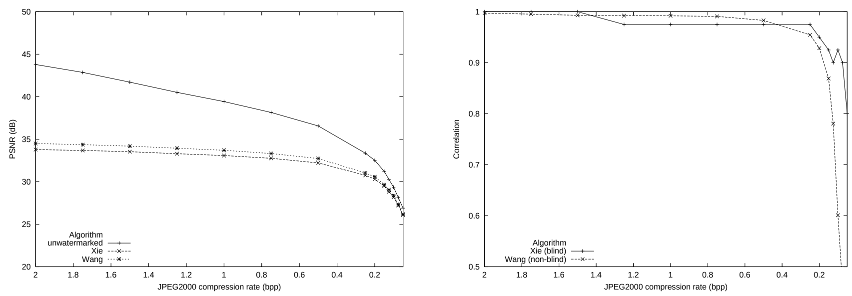

(a) PSNR (b) watermark correlation

Figure 6.11: Illustration of the "distortion gap". The PSNR of the watermarked images is much lower than the compressed version (a). The normalized correlation of the embedded and recovered watermark is shown on the right side (b); at a coding rate of about 0.2 bpp the "distortion gap" closes.

6.4.2 JPEG

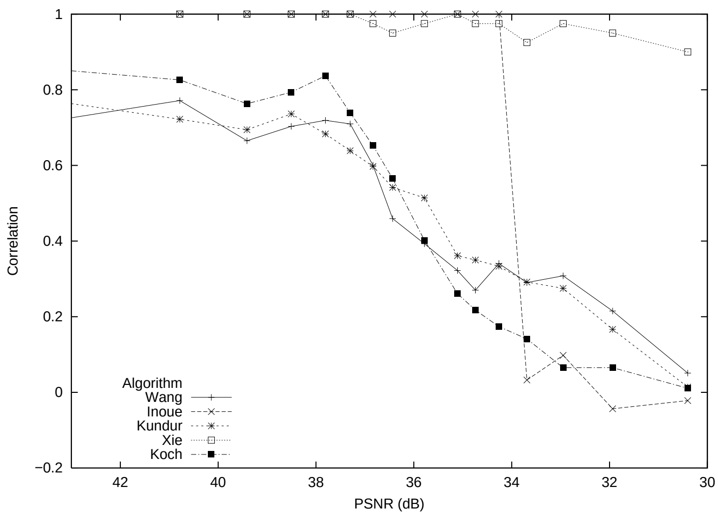

(a) blind algorithms

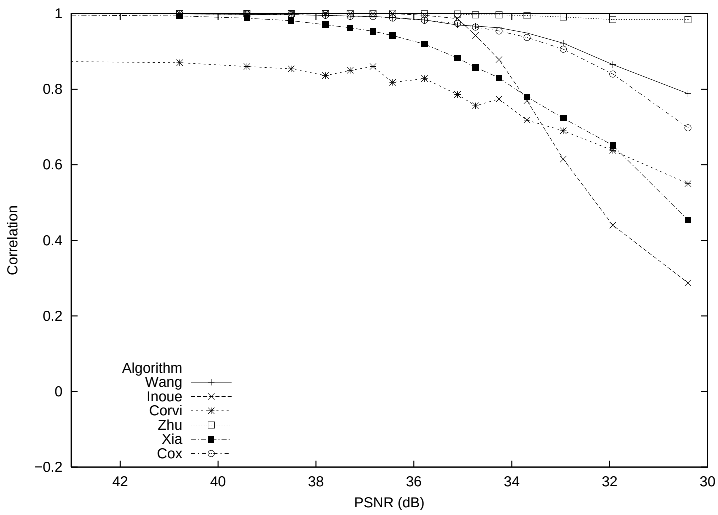

(b) non-blind algorithms

Figure 6.12: Normalized correlation results of the recovered watermark after JPEG image compression. The results for the blind watermarking algorithms are shown in the first row, the second row depicts the results when using a reference image for watermark recovery.

6.4.3 SPIHT

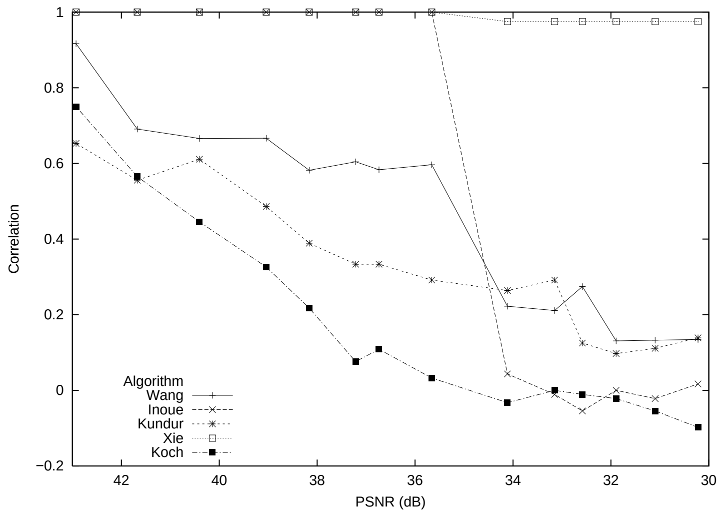

(a) blind algorithms

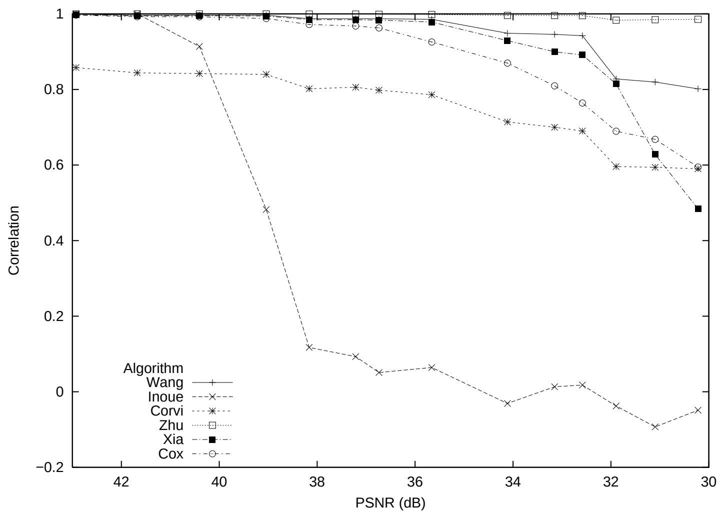

(b) non-blind algorithms

Figure 6.13: Normalized correlation results of the recovered watermark after SPIHT image compression. The results for the blind watermarking algorithms are shown in the first row, the second row depicts the results when using a reference image for watermark recovery.

6.4.4 JPEG2000

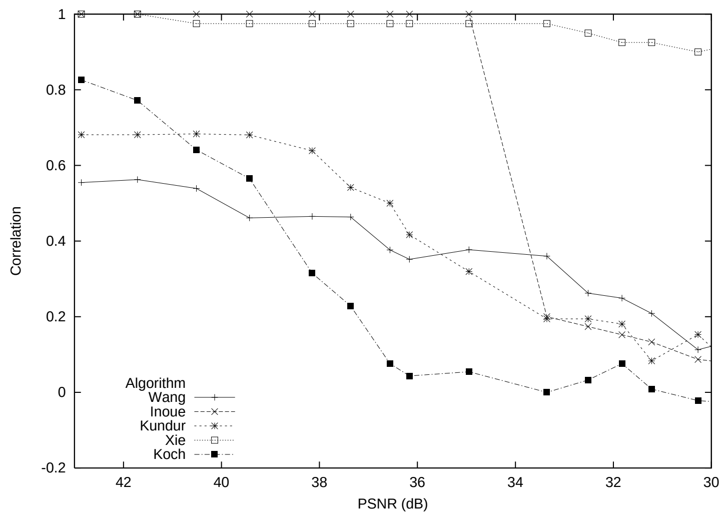

(a) blind algorithms

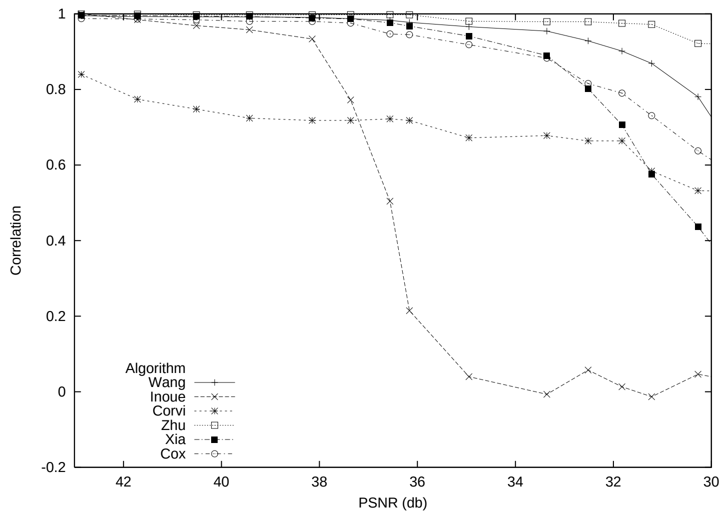

(b) non-blind algorithms

Figure 6.14: Normalized correlation results of the recovered watermark after JPEG2000 image compression. The results for the blind watermarking algorithms are shown in the first row, the second row depicts the results when using a reference image for watermark recovery.

6.5 Image processing

6.5.1 Median filtering

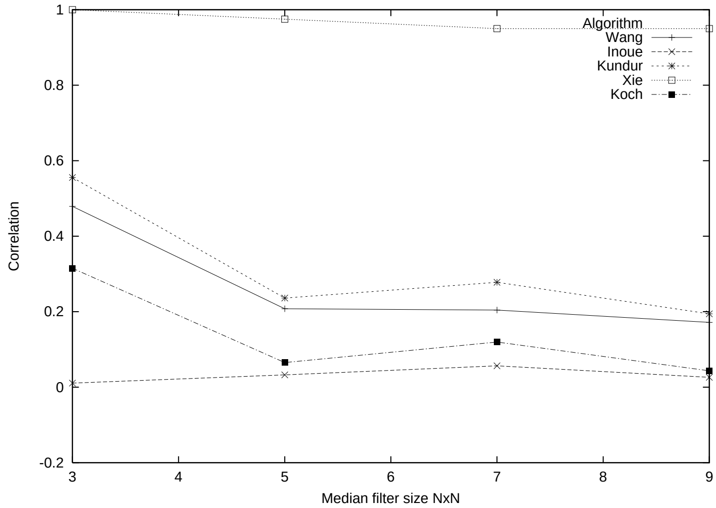

(a) blind algorithms

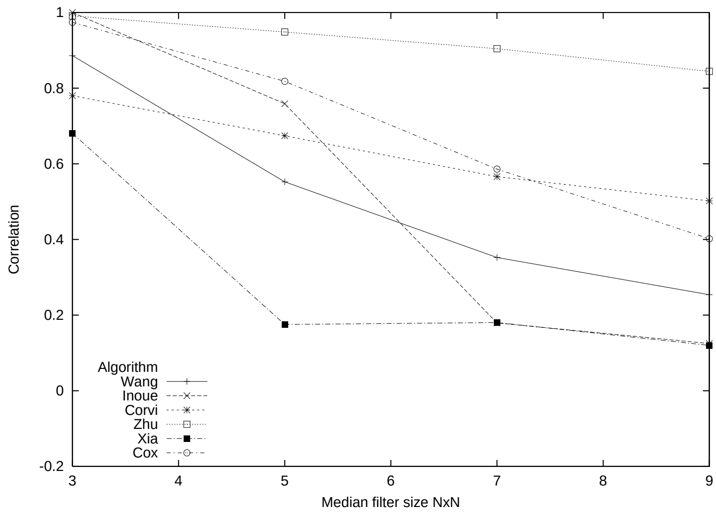

(b) non-blind algorithms

Figure 6.15: Normalized correlation results of the recovered watermark after median filtering attack. The results for the blind watermarking algorithms are shown in the first row, the second row depicts the results when using a reference image for watermark recovery.

6.5.2 Smoothing

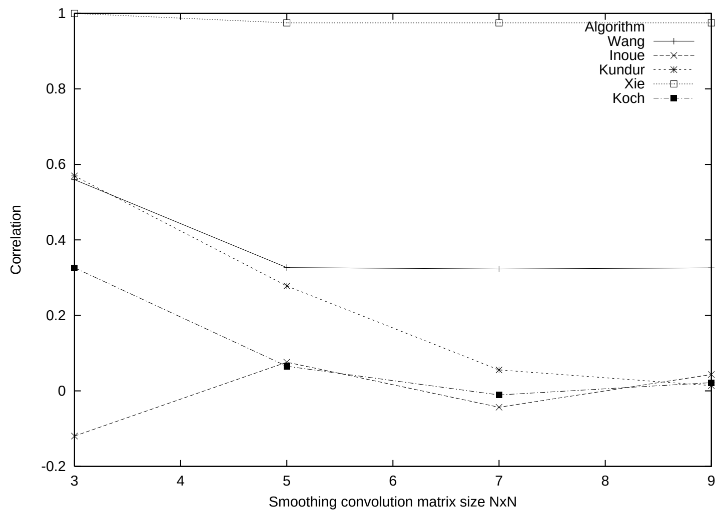

(a) blind algorithms

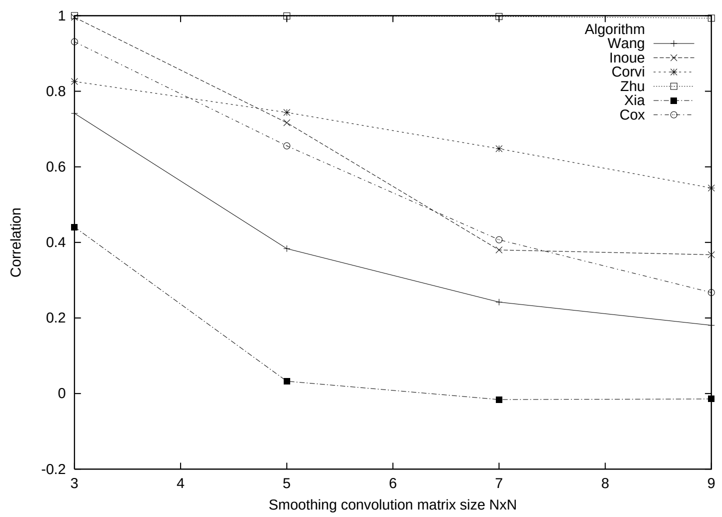

(b) non-blind algorithms

Figure 6.16: Normalized correlation results of the recovered watermark after smoothing attack. The results for the blind watermarking algorithms are shown in the first row, the second row depicts the results when using a reference image for watermark recovery.

6.6 Geometrical transformation

6.6.1 Cropping

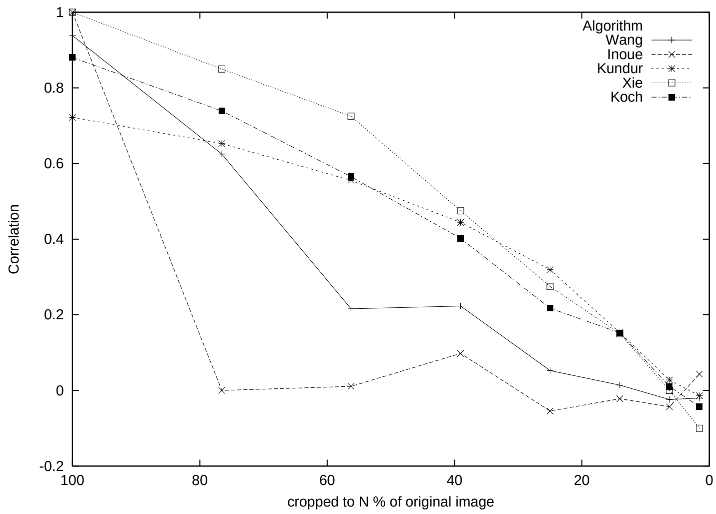

(a) blind algorithms

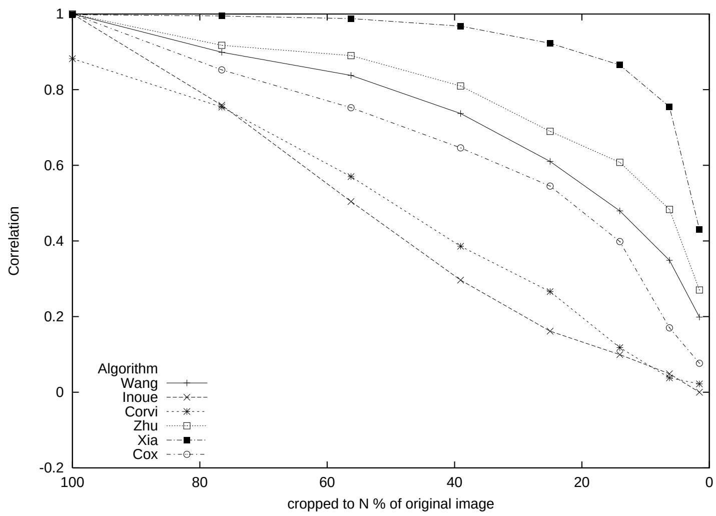

(b) non-blind algorithms

Figure 6.17: Normalized correlation results of the recovered watermark after cropping attack. The results for the blind watermarking algorithms are shown in the first row, the second row depicts the results when using a reference image for watermark recovery.

6.6.2 Down-scaling

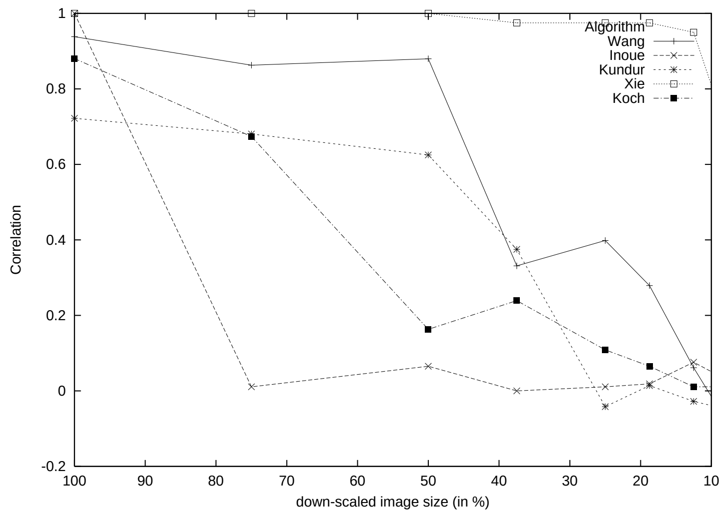

(a) blind algorithms

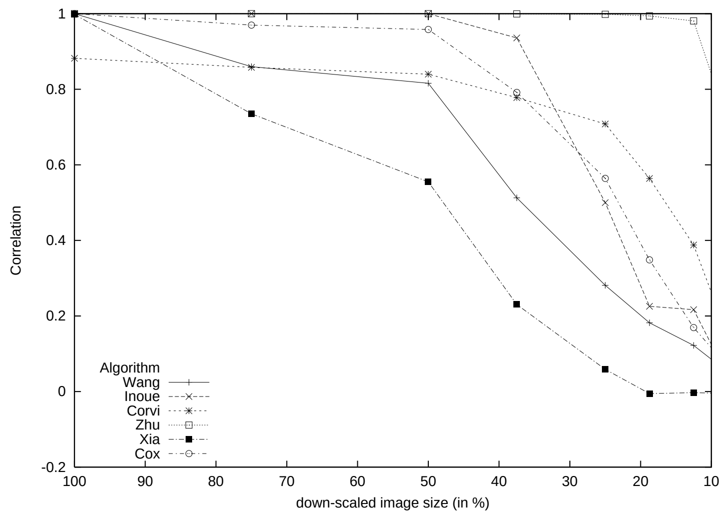

(b) non-blind algorithms

Figure 6.18: Normalized correlation results of the recovered watermark after down-scaling attack. The results for the blind watermarking algorithms are shown in the first row, the second row depicts the results when using a reference image for watermark recovery.

6.7 Image registration

StirMark is probably the most powerful attack against image watermarking. The distortion results from a local non-linear pseudo-random geometrical transformation; see figure 5.1. Most watermarking approaches are also vulnerable to simpler types of geometrical distortion – such as rotation, transposition, scaling and the jitter attack.

With the help of an image registration program, e.g. CREG[2], developed by Loo [144], it is possible to partly revert the distortion caused by StirMark and make watermark detection possible. CREG is based on motion estimation and correction in the complex wavelet domain (CWT) [106, 107] and requires a reference image in order to compute the motion vectors. The reference image can be e.g. a copy of the same image but watermarked with a different key, an undistorted copy of the watermarked image, or the original image.

In figure 6.19 and 6.20, we compare the effect of the StirMark bending attack on the watermarked "Lena" image in terms of PSNR and normalized correlation of the embedded mark. In order to compare the effect of the embedding domain, we depict the results of the spatial- and DCT-domain watermarking methods using the algorithms by Bruyndonckx [20] and Cox [39], Koch [110] and Fridrich [69], respectively, in figure 6.19. The results of the wavelet-domain schemes, using the embedding algorithms of Kim [104], Xia [266], Xie [268] and Zhu [276] are shown in figure 6.20. Moreover, the CREG image registration program was employed to revert the distortion with the help of a reference copy of the watermarked image. It can be seen that some of the wavelet-domain schemes can tolerate modest amounts of StirMark distortion. However, the detection performance of all watermarking schemes, regardless of the embedding domain, can be significantly improved by image registration.

[2]The CREG program is available from `http://www-sigproc.eng.cam.ac.uk/~pl201/watermarking/index.html`.

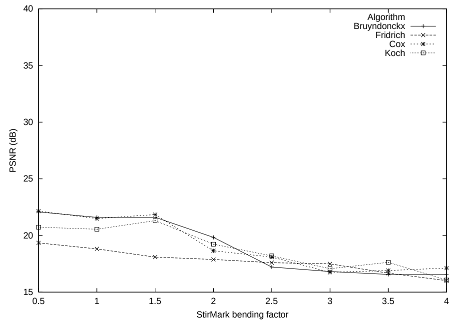

(a) PSNR

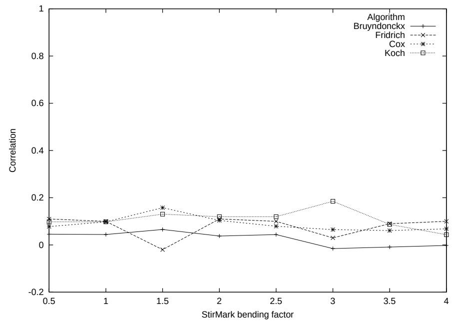

(b) correlation

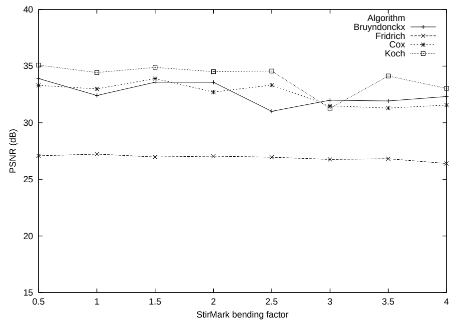

(c) PSNR (registered)

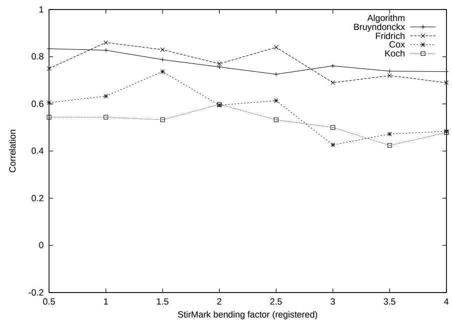

(d) correlation (registered)

Figure 6.19: Image registration of spatial- and DCT domain algorithms. The "Lena" has been watermarked with several algorithms. In the top row, we show the PSNR (a) and the normalized correlation (b) after a StirMark bending attack. The second row illustrates the effect of the CREG registration process; the PSNR (c) as well as the correlation (d) of the watermarked images is improved.

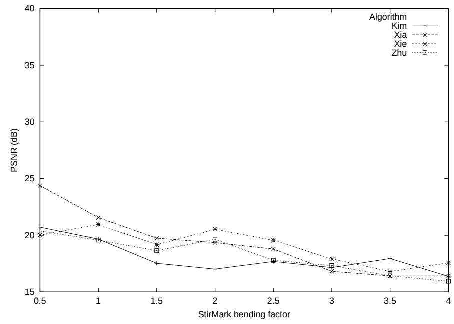

(a) PSNR

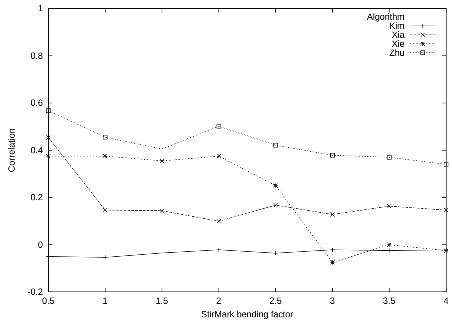

(b) correlation

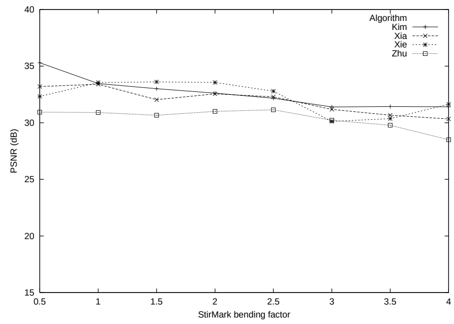

(c) PSNR (registered)

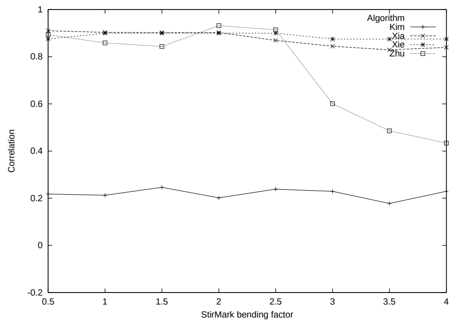

(d) correlation (registered)

Figure 6.20: Image registration of wavelet domain algorithms. The "Lena" has been watermarked with several algorithms. In the top row, we show the PSNR (a) and the normalized correlation (b) after a StirMark bending attack. The second row illustrates the effect of the CREG registration process; the PSNR (c) as well as the correlation (d) of the watermarked images is improved.

Bibliography

[1] JPEG2000 part 1 final committee draft version 1.0. Technical report, ISO/IEC FCD15444-1, March 2000.

[2] Michael Arnold and Sebastian Kanka. MP3 robust audio watermarking. In *DFG V3D2 Watermarking Workshop 1999*, Erlangen, Germany, October 1999.

[3] John Perry Barlow. The economy of ideas. *Wired Magazine*, 2(3), March 1994.

[4] Mauro Barni, Franco Bartolini, Vito Cappellini, Alessandro Lippi, and Alessandro Piva. A DWT-based technique for spatio-frequency masking of digital signatures. In Ping Wah Wong, editor, *Proceedings of the 11th SPIE Annual Symposium, Electronic Imaging '99, Security and Watermarking of Multimedia Contents*, volume 3657, San Jose, CA, USA, January 1999.

[5] Mauro Barni, Franco Bartolini, Vito Cappellini, and Alessandro Piva. Copyright protection of digital images by embedded unperceivable marks. *Image & Vision Computing*, 16(12):897 – 906, August 1998.

[6] Mauro Barni, Franco Bartolini, Vito Cappellini, and Alessandro Piva. A DCT-domain system for robust image watermarking. *Signal Processing, Special Issue on Copyright Protection and Control*, 66(3):357 – 372, May 1998.

[7] Mauro Barni, Franco Bartolini, Vito Cappellini, Alessandro Piva, and F. Rigacci. A m. a. p. identification criterion for DCT-based watermarking. In *Proceedings of the 9th European Signal Processing Conference, EUSIPCO '98*, Island of Rhodes, Greece, 1998.

[8] Mauro Barni, Franco Bartolini, Alessia DeRosa, and Alessandro Piva. Capacity of the watermark channel: how many bits can be hidden within a digital image? In Ping Wah Wong, editor, *Proceedings of the 11th SPIE Annual Symposium, Electronic Imaging '99, Security and Watermarking of Multimedia Contents*, volume 3657, San Jose, CA, USA, January 1999.

[9] P. Bassia and Ioannis Pitas. Robust audio watermarking in the time domain. In *Proceedings of the 9th European Signal Processing Conference, EUSIPCO '98*, Island of Rhodes, Greece, September 1998.

[10] Walter Bender, Daniel Gruhl, Norishige Morimoto, and Anthony Lu. Techniques for data hiding. *IBM Systems Journal*, 35:313 – 336, 1996.

[11] Oliver Benedens. A watermarking system for three-dimensional polygon based models with robustness against mesh simplification. In Ping Wah Wong, editor, *Proceedings of the 11th SPIE Annual Symposium, Electronic Imaging '99, Security and Watermarking of Multimedia Contents*, volume 3657, pages 329 –340, San Jose, CA, USA, January 1999.

[12] Dave Benham, Nasir Memon, Boon-Lock Yeo, and Minerva M. Yeung. Fast watermarking of DCT-based compressed images. In Hamid R. Arabnia, editor, *Proceedings of the International Conference on Image Science, Systems, and Technology, CISST '97*, Las Vegas, USA, 1997.

[13] Anoop K. Bhattacharjya and Hakan Ancin. Data embedding in text for a copier system. In *Proceedings of the 6th IEEE International Conference on Image Processing, ICIP '99*, Kobe, Japan, October 1999.

[14] Jeffrey A. Bloom, Ingemar J. Cox, Ton Kalker, Jean-Paul Linnartz, Matthew L. Miller, and B. Traw. Copy protection for DVD video. *Proceedings of the IEEE Special issue on Identification and Protection of Multimedia Information*, 87(7):1267 – 1276, 1999.

[15] F. M. Boland, Joseph J. K. O'Ruanaidh, and C. Dautzenberg. Watermarking digital images for copyright protection. In *IEE Conference Proceedings on Image Processing and its Applications*, number 410, pages 326 – 330. IEEE, July 1995.

[16] Dan Boneh and James Shaw. Collusion secure fingerprinting for digital data. In Don Coppersmith, editor, *Proceedings of the 15th annual International Cryptology Conference*, volume 963 of *Lecture Notes in Computer Science*, pages 452 – 465, Santa Barbara, CA, USA, August 1995. Springer-Verlag.

[17] Lawrence Boney, Ahmed H. Tewfik, and Khaled N. Hamdy. Digital watermarks for audio signals. In *Proceedings of the IEEE Conference on Multimedia Computing and Systems*, pages 473 – 480, Hiroshima, Japan, 1996.

[18] Adrian G. Bors and Ioannis Pitas. Image watermarking using block site selection and DCT domain constraint. *Optics Express*, 3(12):512, December 1998.

[19] Jack T. Brassil, Steven Low, and Nicholas F. Maxemchuk. Copyright protection for the electronic distribution of text documents. *Proceedings of the IEEE, Special Issue on Protection of Multimedia Content*, 87(2):1181 – 1196, July 1999.

[20] O. Bruyndonckx, Jean-Jacques Quisquater, and Benoit M. Macq. Spatial method for copyright labeling of digital images. In *Proceedings of the IEEE International Workshop on Nonlinear Signal and Image Processing*, pages 456 – 459, Marmaras, Greece, 1995.

[21] Jong Jin Chae. *Robust Techniques for Data Hiding in Images and Video*. PhD thesis, Department for Electrical and Computer Engineering, University of California, Santa Barbara, CA, USA, 1999.

[22] Jong Jin Chae and B. S. Manjunath. Extracting hidden data without knowing host source. In Ping Wah Wong, editor, *Proceedings of the 11th SPIE Annual Symposium, Electronic Imaging '99, Security and Watermarking of Multimedia Contents*, volume 3657, San Jose, CA, USA, January 1999.

[23] Jong Jin Chae and B. S. Manjunath. A technique for image data hiding and reconstruction without host image. In Ping Wah Wong, editor, *Proceedings of the 11th SPIE Annual Symposium, Electronic Imaging '99, Security and Watermarking of Multimedia Contents*, volume 3657, pages 386 – 396, San Jose, CA, USA, January 1999.

[24] Jong Jin Chae, Debargha Mukherjee, and B. S. Manjunath. Color image embedding using multidimensional lattice structures. In *Proceedings of the IEEE International Conference on Image Processing, ICIP '98*, pages 460 – 464, Chicago, IL, USA, October 1998.

[25] Jong Jin Chae, Debargha Mukherjee, and B. S. Manjunath. Robust data hiding technique using multidimensional lattices. In *Proceedings of the IEEE Advances in Digital Libraries Conference*, pages 319 – 326, Santa Barbara, CA, USA, April 1998.

[26] Jong Jin Chae, Debargha Mukherjee, and B. S. Manjunath. A robust embedded data from wavelet coefficients. In *Proceedings of SPIE, Electronic Imaging, Storage and Retrieval for Image and Video Database*, volume 3312, pages 308 – 317, San Jose, CA, USA, January 1998.

[27] Maryline Charrier, Diego Santa Cruz, and Mathias Larsson. JPEG2000, the next millennium compression standard for still images. In *Proceedings of the IEEE International Conference on Multimedia & Computing Systems, ICMCS '99*, volume 1, pages 131 – 132, Florence, Italy, June 1999.

[28] Brian Chen and Gregory W. Wornell. Digital watermarking and information embedding using dither modulation. In *Proceedings of the IEEE Workshop on Multimedia Signal Processing*, pages 273 – 278, Los Angeles, CA, USA, December 1998. IEEE.

[29] Brian Chen and Gregory W. Wornell. Dither modulation: A new approach to digital watermarking and information embedding. In Ping Wah Wong, editor, *Proceedings of the 11th SPIE Annual Symposium, Electronic Imaging '99, Security and Watermarking of Multimedia Contents*, volume 3657, pages 342 – 353, San Jose, CA, USA, January 1999.

[30] Qing Chen and Yanping Wang. Robust image watermarking based on embedded zerotree wavelet. unpublished.

[31] Chee-Jung Henry Chu and Anthony Wayne Wiltz. Luminance channel modulated watermarking of digital images. In *Proceedings of the SPIE Wavelet Applications Conference*, pages 437 – 445, Orlando, FL, USA, April 1999.

[32] Christian Collberg and Clark Thomborson. On the limits of software watermarking. Technical report, Department of Computer Science, University of Auckland, New Zealand, August 1998.

[33] Christian Collberg and Clark Thomborson. Software watermarking: Models and dynamic embedding. In *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL '99*, pages 311 – 324, San Antonio, TX, USA, January 1999.

[34] Marco Corvi and Gianluca Nicchiotti. Wavelet-based image watermarking for copyright protection. In *Scandinavian Conference on Image Analysis, SCIA '97*, Lappeenranta, Finland, June 1997.

[35] Max H. M. Costa. Writing on dirty paper. *IEEE Transactions on Information Theory*, 29(3):439 – 441, May 1983.

[36] Ingemar J. Cox, Joe Kilian, Tom Leighton, and Talal G. Shamoon. Secure spread spectrum watermarking for multimedia. Technical report, NEC Research Institure, Princeton, USA, October 1995.

[37] Ingemar J. Cox, Joe Kilian, Tom Leighton, and Talal G. Shamoon. A secure, robust watermark for multimedia. In Ross Anderson, editor, *Information hiding: first international workshop*, volume 1174 of *Lecture Notes in Computer Science*, pages 183 – 206, Cambridge, UK, 1996. Springer Verlag, Berlin, Germany.

[38] Ingemar J. Cox, Joe Kilian, Tom Leighton, and Talal G. Shamoon. Secure spread spectrum watermarking for images, audio and video. In *Proceedings of the IEEE International Conference on Image Processing, ICIP '96*, pages 243 – 246, Lausanne, Switzerland, September 1996. IEEE Press.

[39] Ingemar J. Cox, Joe Kilian, Tom Leighton, and Talal G. Shamoon. Secure spread spectrum watermarking for multimedia. In *Proceedings of the IEEE International Conference on Image Processing, ICIP '97*, volume 6, pages 1673 – 1687, Santa Barbara, California, USA, October 1997.

[40] Ingemar J. Cox, Matthew L. Miller, and Andrew L. McKellips. Watermarking as communications with side information. *Proceedings of the IEEE, Special Issue on Protection of Multimedia Content*, 87(2):1127 – 1141, July 1999.

[41] Scott Craver. On public-key steganography in the presence of an active warden. In David Aucsmith, editor, *Information hiding: second international workshop*, volume 1525 of *Lecture notes in computer science*, Portland, OR, USA, April 1998. Springer Verlag, Berlin, Germany.

[42] Scott Craver, Nasir Memon, Boon-Lock Yeo, and Minerva M. Yeung. Can invisible watermarks resolve rightful ownerships? Technical Report 20509, IBM Research Report, July 1996.

[43] Scott Craver, Nasir Memon, Boon-Lock Yeo, and Minerva M. Yeung. On the invertibility of invisible watermarking techniques. In *Proceedings of the IEEE International Conference on Image Processing, ICIP '97*, volume 1, page 540, Santa Barbara, California, USA, October 1997.

[44] Scott Craver, Nasir Memon, Boon-Lock Yeo, and Minerva M. Yeung. Resolving rightful ownership with invisible watermarking techniques: Limitations, attacks and implications. *IEEE Journal, Selected Areas of Communications*, 16(4):573 – 586, May 1998.

[45] Scott Daly, Wenjun Zeng, Jin Li, and Shawmin Lei. Visual masking in wavelet compression for JPEG2000. In Ping Wah Wong, editor, *Proceedings of IS&T/SPIE's 12th Annual Symposium, Electronic Imaging 2000: Security and Watermarking of Multimedia Content II*, volume 3971, San Jose, CA, USA, January 2000.

[46] Ingrid Daubechies. Orthonormal bases of compactly supported wavelets. *Communication in Pure and Applied Mathematics*, (41):909 – 996, 1988.

[47] Ingrid Daubechies. *Ten lectures on wavelets*. SIAM Press, Philadelphia, PA, USA, 1992.

[48] Geoffrey M. Davis and Aria Nosratinia. Wavelet-based image coding: An overview. *Applied and Computational Control, Signals, and Circuits*, 1(1), 1998.

[49] Franck Davoine. Triangular meshes: a solution to resist to geometric distortions based watermark-removal softwares. In *Proceedings of the 10th European Signal Processing Conference, EUSIPCO '00*, Tampere, Finland, September 2000.

[50] Frédéric Deguillaume, Gabriella Csurka, and Thierry Pun. Countermeasures for unintentional and intentional video watermarking attacks. In Ping Wah Wong, editor, *Proceedings of IS&T/SPIE's 12th Annual Symposium, Electronic Imaging 2000: Security and Watermarking of Multimedia Content II*, volume 3971, San Jose, CA, USA, January 2000.

[51] Philippe Desarte, Benoit Macq, and Dirk T. M. Slock. Signal-adapted multiresolution transform for image coding. *IEEE Transactions on Information Theory*, 38(2), March 1992.

[52] Jana Dittmann, Alexander Behr, Mark Stabenau, Peter Schmitt, Jörg Schwenk, and Johannes Ueberberg. Combining digital watermarks and collusion secure fingerprints for digital images. In Ping Wah Wong, editor, *Proceedings of the 11th SPIE Annual Symposium, Electronic Imaging '99, Security and Watermarking of Multimedia Contents*, volume 3657, San Jose, CA, USA, January 1999.

[53] Jana Dittmann, Mark Stabenau, and Ralf Steinmetz. Robust MPEG video watermarking technologies. In *Multimedia and Security Workshop at the 6th ACM International Multimedia Conference*, pages 113 – 122, Bristol, England, 1998.

[54] Stefan Droste. New results on visual cryptography. In Ross Anderson, editor, *Information hiding: first international workshop*, volume 1174 of *Lecture Notes in Computer Science*, pages 401 – 412, Cambridge, UK, 1996. Springer Verlag, Berlin, Germany.

[55] Rakesh Dugad, Krishna Ratakonda, and Narendra Ahuja. A new wavelet-based scheme for watermarking images. In *Proceedings of the IEEE International Conference on Image Processing, ICIP '98*, Chicago, IL, USA, October 1998.

[56] Jean-Luc Dugelay and Fabien A. P. Petitcolas. Possible counter-attacks against random geometric distortions. In Ping Wah Wong, editor, *Proceedings of IS&T/SPIE's 12th Annual Symposium, Electronic Imaging 2000: Security and Watermarking of Multimedia Content II*, volume 3971, San Jose, CA, USA, January 2000.

[57] Jean-Luc Dugelay and Stephane Roche. Fractal transform based large digital watermark embedding and robust full blind extraction. In *Proceedings of the IEEE International Conference on Multimedia & Computing Systems, ICMCS '99*, volume 2, pages 1003 – 1004, Florence, Italy, June 1999.

[58] Esther Dyson. Intellectual value. *Wired Magazine*, 3(7), July 1995.

[59] Joachim J. Eggers and Bernd Girod. Watermark detection after quantization attacks. In *Proceedings of the 3rd Information Hiding Workshop '99*, volume 1768, pages 172 – 186, Dresden, Germany, September 1999. Springer.

[60] Joachim J. Eggers and Bernd Girod. Blind watermarking applied to image authentication. In *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing, ICASSP '01*, Salt Lake City, UT, USA, May 2001.

[61] Joachim J. Eggers, Wolf-D. Ihlenfeldt, and Bernd Girod. Digital watermarking of chemical structure sets. In *Proceedings of the 4th Information Hiding Workshop '01*, Portland, OR, USA, April 2001.

[62] Joachim J. Eggers, Jonathan K. Su, and Bernd Girod. Asymmetric watermarking schemes. In *Tagungsband der Gesellschaft für Informatik zur 30. Jahrestagung*, pages 124 – 133, Berlin, Germany, September 2000.

[63] Joachim J. Eggers, Jonathan K. Su, and Bernd Girod. A blind watermarking scheme based on structured codebooks. In *IEE Colloquium: Secure images and image authentication*, London, UK, April 2000.

[64] Joachim J. Eggers, Jonathan K. Su, and Bernd Girod. Public key watermarking by eigenvectors of linear transforms. In *Proceedings of the 10th European Signal Processing Conference, EUSIPCO '00*, Tampere, Finland, September 2000.

[65] Masataka Ejima and Akio Miyazaki. A wavelet-based watermarking for digital images and video. In *Proceedings of the IEEE International Conference on Image Processing, ICIP '00*, Vancouver, Canada, September 2000.

[66] Masataka Ejima and Akio Miyazaki. A wavelet-based watermarking for digital images and video. *IEICE Transactions*, 83(3):532 – 540, March 2000.

[67] Masataka Ejima, Akio Miyazaki, and Taku Saito. Digital watermark based on the dyadic wavelet transform and its robustness on image compressing. In *1998 International Technical Conference on Circuits/Systems, Computers and Communications*, Sokcho, Korea, July 1998.

[68] J. Mark Ettinger. Steganalysis and game equilibria. In David Aucsmith, editor, *Information hiding: second international workshop*, volume 1525 of *Lecture notes in computer science*, Portland, OR, USA, April 1998. Springer Verlag, Berlin, Germany.

[69] Jiri Fridrich. Combining low-frequency and spread spectrum watermarking. In *Proceedings of the SPIE Symposium on Optical Science, Engineering and Instrumentation*, San Diego, USA, July 1998.

[70] Jiri Fridrich. Key-dependent random image transforms and their applications in image watermarking. In *Proceedings of the 1999 International Conference on Imaging Science, Systems, and Technology, CISST '99*, pages 237 – 243, Las Vegas, NV, USA, June 1999.

[71] Jiri Fridrich, Arnold C. Baldoza, and Richard J. Simard. Robust digital watermarking based on key-dependent basis functions. In David Aucsmith, editor, *Information hiding: second international workshop*, volume 1525 of *Lecture notes in computer science*, pages 143 – 157, Portland, OR, USA, April 1998. Springer Verlag, Berlin, Germany.

[72] Bernd Girod. The information theoretical significance of spatial and temporal masking in video signals. In *Proceedings of the SPIE Symposium on Electronic Imaging, Conference on Human Vision, Visual Processing and Digital Display*, volume 1077, pages 178 – 187, Los Angeles, CA, USA, 1989.

[73] Henry M. Gladney, Frederick C. Mintzer, and Fabio Schiattarella. Safeguarding digital library contents and users: Digital images of treasured antiquities. *D-Lib Magazine*, July 1997.

[74] Henry M. Gladney, Frederick C. Mintzer, Fabio Schiattarella, Julien Besces, and Martin Treu. Digital access to antiquities. *Communications of the ACM*, 41(4):49 – 57, April 1998.

[75] Frank Hartung and Bernd Girod. Digital watermarking of raw and compressed video. In Naohisa Ohta, editor, *Proceedings of the SPIE Digital Compression Technologies and Systems for Video Communications Conference*, volume 2952, pages 205 – 213, October 1996.

[76] Frank Hartung and Martin Kutter. Multimedia watermarking techniques. In *Proceedings of the IEEE, Special Issue on Protection of Multimedia Content*, volume 87, pages 1079 – 1107, July 1999.

[77] Frank Hartung, Jonathan K. Su, and Bernd Girod. Spread spectrum watermarking: Malicious attacks and counter-attacks. In Ping Wah Wong, editor, *Proceedings of the 11th SPIE Annual Symposium, Electronic Imaging '99, Security and Watermarking of Multimedia Contents*, volume 3657, San Jose, CA, USA, January 1999.

[78] Juan Ramon Hernández and Fernando Pérez-González. Statistical analysis of watermarking schemes for copyright protection of images. *Proceedings of the IEEE, Special Issue on Protection of Multimedia Content*, 87(2):1142 – 1165, July 1999.

[79] Alexander Herrigel, Joseph J. K. O'Ruanaidh, Holger Petersen, Shelby Pereira, and Thierry Pun. Secure copyright protection techniques for digital images. In David Aucsmith, editor, *Information hiding: second international workshop*, volume 1525 of *Lecture notes in computer science*, pages 169 – 190, Portland, OR, USA, April 1998. Springer Verlag, Berlin, Germany.

[80] Zenon Hrytskiv, Sviatoslav Voloshynovskiy, and Y. B. Rytsar. Cryptography and steganography of video information in modern communications. In *Proceedings of the 3rd International Conference on Telecommunications in Modern Satelite, Cable and Broadcasting Services TELSIKS '97*, volume 1, pages 164 – 167, Nis, Yugoslavia, 1997.

[81] Chiou-Ting Hsu and Ja-Ling Wu. Multiresolution watermarking for digital images. *IEEE Transactions on Circuits and Systems II*, 45(8):1097 – 1101, August 1998.

[82] Wolf-D. Ihlenfeldt and Joachim J. Eggers. Watermarking of chemical structure sets. Technical report, Computer Chemistry Center, University of Nuremberg-Erlangen, Germany, April 1999.

[83] Hisashi Inoue, Akio Miyazaki, Takashi Araki, and Takashi Katsura. A digital watermark method using the wavelet transform for video data. *IEICE Transactions*, 83(1):90 – 96, January 2000.

[84] Hisashi Inoue, Akio Miyazaki, and Takashi Katsura. A digital watermark for image signals using a controlled quantization. In *1998 International Technical Conference on Circuits/Systems, Computers and Communications*, Sokcho, Korea, 1998.

[85] Hisashi Inoue, Akio Miyazaki, and Takashi Katsura. An image watermarking method based on the wavelet transform. In *Proceedings of the 6th IEEE International Conference on Image Processing, ICIP '99*, page 375, Kobe, Japan, October 1999.

[86] Hisashi Inoue, Akio Miyazaki, and Takashi Katsura. Wavelet-based watermarking for tamper proofing of still images. In *Proceedings of the IEEE International Conference on Image Processing, ICIP '00*, Vancouver, Canada, September 2000.

[87] Hisashi Inoue, Akio Miyazaki, Akihiro Yamamoto, and Takashi Katsura. A digital watermark based on the wavelet transform and its robustness on image compression. In *Proceedings of the IEEE International Conference on Image Processing, ICIP '98*, Chicago, IL, USA, 1998.

[88] Hisashi Inoue, Akio Miyazaki, Akihiro Yamamoto, and Takashi Katsura. A digital watermark technique based on the wavelet transform and its robustness on image compression and transformation. *IEICE Transactions, Special Section on Cryptography and Information Security*, 82(1):11 – 20, January 1999.

[89] Ashoka Jayawardena, Bob Murison, and Patrick Lenders. High bit rate wavelet domain digital watermarking of images and compression tolerance. In *Fourth Australasian Document Computing Symposium*, Australia, December 1999.

[90] Ashoka Jayawardena, Bob Murison, and Patrick Lenders. Embedding multiresolution binary images into multiresolution watermark channels in wavelet domain. In *Proceedings of the IEEE ICASSP 2000*, Istanbul, Turkey, June 2000.

[91] Brigitte Jellinek. Invisible watermarking of digital images for copyright protection. Master's thesis, Department of Scientific Computing, University of Salzburg, Austria, January 2000.

[92] N. Kaewkamnerd and K. R. Rao. Wavelet based image adaptive watermarking scheme. *Electronic Letters*, 36(4):312 – 313, February 2000.

[93] Ton Kalker. A security risk for publicly available watermark detectors. In *Proceedings of the 18th Symposium on Information Theory in the Benelux*, Veldhoven, The Netherlands, 1998.

[94] Ton Kalker. Watermark estimation through detector observations. In *Proceedings of the Benelux Signal Processing Symposium*, Leuven, Belgium, 1998.

[95] Ton Kalker. Digital video watermarking for DVD copy protection. In *DFG V3D2 Watermarking Workshop 1999*, Erlangen, Germany, October 1999.

[96] Ton Kalker, Geert Depovere, Jaap Haitsma, and Maurice Maes. A video watermarking system for broadcast monitoring. In Ping Wah Wong, editor, *Proceedings of the 11th SPIE Annual Symposium, Electronic Imaging '99, Security and Watermarking of Multimedia Contents*, volume 3657, pages 103 – 112, San Jose, CA, USA, January 1999.

[97] Ton Kalker, Jean-Paul Linnartz, Geert Depovere, and Maurice Maes. On the reliability of detecting electronic watermarks in digital images. In *Proceedings of the 9th European Signal Processing Conference, EUSIPCO '98*, pages 13 – 16, Island of Rhodes, Greece, September 1998.

[98] Ton Kalker, Jean-Paul Linnartz, and Marten van Dijk. Watermark estimation through detector analysis. In *Proceedings of the IEEE International Conference on Image Processing, ICIP '98*, Chicago, IL, USA, October 1998.

[99] Satoshi Kanai, Hiroaki Date, and Takeshi Kishinami. Digital watermarking for 3d polygons using multiresolution wavelet decomposition. In *Proceedings of the Sixth IFIP WG 5.2 International Workshop on Geometric Modeling: Fundamentals and Applications*, pages 296 – 307, Tokyo, Japan, December 1998.

[100] Auguste Kerckhoff. La cryptographie militaire. *Journal des sciences militaires*, 9:5 – 38, January 1883.

[101] Sanjeev Khanna and Francis Zane. Watermarking maps: Hiding information in structured data. In *Proceedings of the 11th Annual ACM-SIAM Symposium on Discrete Algorithms*, San Francisco, CA, USA, January 2000.

[102] Joe Kilian, Tom Leighton, Lesley R. Matheson, Talal G. Shamoon, Robert E. Tarjan, and Francis Zane. Resistance of digital watermarks to collusive attacks. In *Proceedings of 1998 IEEE International Symposium on Information Theory, ISIT '98*, page 271, Cambridge, MA, USA, August 1998.

[103] Hongseok Kim. Stochastic model based audio watermark and whitening filter for improved detection. In *Proceedings of the IEEE ICASSP 2000*, Istanbul, Turkey, June 2000.

[104] Jong Ryul Kim and Young Shik Moon. A robust wavelet-based digital watermark using level-adaptive thresholding. In *Proceedings of the 6th IEEE International Conference on Image Processing, ICIP '99*, page 202, Kobe, Japan, October 1999.

[105] Y.-S. Kim, O.-H. Kwon, and R.-H. Park. Wavelet based watermarking method for digital images using the human visual system. *Electronic Letters*, 35(6):466 – 467, June 1999.

[106] Nick G. Kingsbury. The dual-tree complex wavelet transform: a new technique for shift invariance and directional filters. In *Proceedings of the IEEE Digital Signal Processing Workshop, DSP '98*, Bryce Canyon, USA, August 1998.

[107] Nick G. Kingsbury. Image processing with complex wavelets. In *Phil. Trans. of the Royal Society London, Discussion meeting on "Wavelets: the key to intermittent information?"*, London, UK, February 1999.

[108] Nick G. Kingsbury and D. W. Redmill. Image and video wavelet coding for noisy channels. In *Proceedings of the SPIE AeroSense Conference '97*, Orlando, FL, USA, April 1997.

[109] Donald E. Knuth. *The Art of Computer Programming, Seminumerical Algorithms*, volume 2. Addison-Wesley, edition.1228 edition, October 1998.

[110] Eckhard Koch and Jian Zhao. Towards robust and hidden image copyright labeling. In *Proceedings of the IEEE International Workshop on Nonlinear Signal and Image Processing*, pages 452 – 455, Marmaras, Greece, June 1995.

[111] Deepa Kundur. Improved digital watermarking through diversity and attack characterization. In *Proceedings of the ACM Workshop on Multimedia Security '99*, pages 53 – 58, Orlando, FL, USA, October 1999.

[112] Deepa Kundur. *Multiresolution Digital Watermarking: Algorithms and Implications for Multimedia Signals*. PhD thesis, Dept. of Electrical & Computer Engineering, University of Toronto, Canada, August 1999.

[113] Deepa Kundur. Energy allocation for high-capacity watermarking in the presence of compression. In *Proceedings of the IEEE International Conference on Image Processing, ICIP '00*, Vancouver, Canada, September 2000.

[114] Deepa Kundur and Dimitrios Hatzinakos. A robust digital image watermarking method using wavelet-based fusion. In *Proceedings of the IEEE International Conference on Image Processing, ICIP '97*, volume 1, pages 544 – 547, Santa Barbara, California, USA, October 1997.

[115] Deepa Kundur and Dimitrios Hatzinakos. Digital watermarking using multiresolution wavelet decomposition. In *Proceedings of IEEE ICASSP '98*, volume 5, pages 2969 – 2972, Seattle, WA, USA, May 1998.

[116] Deepa Kundur and Dimitrios Hatzinakos. Improved robust watermarking through attack characterization. *Optics Express*, 3(12):485, December 1998.

[117] Deepa Kundur and Dimitrios Hatzinakos. Semi-blind image restoration based on telltale watermarking. In *Proceedings of the 32nd Asilomar Conference on Signals, Systems, and Computers*, 1998.

[118] Deepa Kundur and Dimitrios Hatzinakos. Towards a telltale watermarking technique for tamper-proofing. In *Proceedings of the IEEE International Conference on Image Processing, ICIP '98*, volume 2, pages 409 – 413, Chicago, IL, USA, October 1998.

[119] Deepa Kundur and Dimitrios Hatzinakos. Attack characterization for effective watermarking. In *Proceedings of the 6th IEEE International Conference on Image Processing, ICIP '99*, Kobe, Japan, October 1999.

[120] Deepa Kundur and Dimitrios Hatzinakos. Digital watermarking for telltale tamper-proofing and authentication. In *Proceedings of the IEEE: Special Issue on Identification and Protection of Multimedia Information*, volume 87, pages 1167 – 1180, July 1999.

[121] Deepa Kundur and Dimitrios Hatzinakos. Mismatching perceptual models for effective watermarking in the presence of compression. In *Proceedings of the SPIE Conference on Multimedia Systems and Applications II*, volume 3845, Boston, MA, USA, September 1999.

[122] Martin Kutter. *Digital Image Watermarking: Hiding Information in Images*. PhD thesis, EPFL, Lausanne, Switzerland, 1999.

[123] Martin Kutter. Watermarking resisting to translation, rotation, and scaling. In *Proceedings of SPIE: Multimedia systems and applications*, volume 3528, pages 423 – 431, 1999.

[124] Martin Kutter, Sushil K. Bhattacharjee, and Touradj Ebrahimi. Towards second generation watermarking schemes. In *Proceedings of the 6th IEEE International Conference on Image Processing, ICIP '99*, volume 1, pages 320 – 323, Kobe, Japan, October 1999.

[125] Martin Kutter, Frédéric Jordan, and Frank Bossen. Digital signature of color images using amplitude modulation. In Ishwar K. Sethi, editor, *Proceedings of the SPIE Conference on Storage and Retrieval for Image and Video Databases*, volume 2952, pages 518 – 526, San Jose, USA, 1997.

[126] Martin Kutter and Franck Leprévost. Symbiose von kryptographie und digitalen Wasserzeichen: effizienter Schutz des Urheberrechtes digitaler Medien. In *Tagungsverband des 6. Deutschen IT-Sicherheitskongreß, Bundesamt für Sicherheit in der Informationstechnik*, pages 479 – 484, May 1999.

[127] Martin Kutter and Fabien A. P. Petitcolas. A fair benchmark for image watermarking systems. In Ping Wah Wong, editor, *Proceedings of the 11th SPIE Annual Symposium, Electronic Imaging '99, Security and Watermarking of Multimedia Contents*, volume 3657, pages 226 – 239, San Jose, CA, USA, January 1999.

[128] Martin Kutter, Sviatoslav Voloshynovskiy, and Alexander Herrigel. Watermark copy attack. In Ping Wah Wong, editor, *Proceedings of IS&T/SPIE's 12th Annual Symposium, Electronic Imaging 2000: Security and Watermarking of Multimedia Content II*, volume 3971, San Jose, CA, USA, January 2000.

[129] Gerrit C. Langelaar, Reginald L. Lagendijk, and Jan Biemond. Real-time labeling methods for MPEG compressed video. In *Proceedings of the 18th Symposium on Information Theory in the Benelux*, Veldhoven, The Netherlands, February 1997.

[130] Gerrit C. Langelaar, Reginald L. Lagendijk, and Jan Biemond. Removing spatial spread spectrum watermarks by non-linear filtering. In *Proceedings of the 9th European Signal Processing Conference, EUSIPCO '98*, Island of Rhodes, Greece, September 1998.

[131] Jacques Levy-Vehel and Anne Manoury. Wavelet packet based digital watermarking. In *Proceedings of the 15th International Conference on Pattern Recognition*, Barcelona, Spain, September 2000.

[132] A. S. Lewis and G. Knowles. Image compression using the 2-d wavelet transform. *IEEE Transactions on Image Processing*, 1:244 – 250, April 1992.

[133] Xin Li and Hong Heather Yu. Transparent and robust audio data hiding in subband domain. In *Proceedings of the IEEE International Conference on Information Technology: Coding and Computing*, Las Vegas, NV, USA, March 2000.

[134] J. Liang, P. Xu, and T. D. Tran. A universal robust low frequency watermarking scheme. *submitted to IEEE Transactions on Image Processing*, May 2000.

[135] Vinicius Licks, R. Jordan, D. F. G. Azevedo, J. S. Correa, P. R. G. Franco, and R. D. R Fagundes. Circular watermark robust to geometric attacks. In *Proceedings of the IEEE International Symposium on Information Theory and Its Applications*, Hawai, USA, 2000.

[136] Ching-Yung Lin and Shih-Fu Chang. A watermark-based robust image authentication method using wavelets. Technical Report ADVENT, Department of Electrical Engineering, Columbia University, NY, USA, April 1998.

[137] Jean-Paul Linnartz, Ton Kalker, and Geert Depovere. Modelling the false alarm and missed detection rate for electronic watermarks. In David Aucsmith, editor, *Information hiding: second international workshop*, volume 1525 of *Lecture notes in computer science*, pages 329 – 343, Portland, OR, USA, April 1998. Springer Verlag, Berlin, Germany.

[138] Jean-Paul Linnartz, Ton Kalker, and Jaap Haitsma. Detecting electronic watermarks in digital video. In *Proceedings of ICASSP '99*, volume 4, pages 2071 – 2074, Phoenix, AZ, USA, March 1999.

[139] Jean-Paul Linnartz and J. C. Talstra. MPEG PTY-marks: Cheap detection of embedded copyright data in DVD-video. In *Fifth European Symposium on Research in Computer Security ESORICS '98*, volume 1485, pages 221 – 240, September 1998.

[140] Jean-Paul Linnartz and Marten van Dijk. Analysis of the sensitivity attack against electronic watermarks in images. In David Aucsmith, editor, *Information hiding: second international workshop*, volume 1525 of *Lecture notes in computer science*, pages 258 – 272, Portland, OR, USA, April 1998. Springer Verlag, Berlin, Germany.

[141] Yong Liu, Jonathan Mant, Edward Wong, and Steven Low. Marking and detection of text documents using transform-domain techniques. In Ping Wah Wong, editor, *Proceedings of the 11th SPIE Annual Symposium, Electronic Imaging '99, Security and Watermarking of Multimedia Contents*, volume 3657, San Jose, CA, USA, January 1999.

[142] Patrick Loo and Nick G. Kingsbury. Digital watermarking using complex wavelets (trimmed version). Technical report, Department of Engineering, Cambridge University, UK, August 1999.

[143] Patrick Loo and Nick G. Kingsbury. Watermarking using complex wavelets with resistance to geometric distortion. In *Proceedings of the 10th European Signal Processing Conference, EUSIPCO '00*, Tampere, Finland, September 2000.

[144] Patrick Loo and Nick G. Kingsbury. Motion estimation based registration of geometrically distorted images for watermark recovery. In *Proceedings of SPIE, Security and Watermarking of Multimedia Contents*, volume 4314, San Jose, CA, USA, January 2001.

[145] Chun-Shien Lu, Yih-Feng Chen, Hong-Yuan Mark Liao, and Chiou-Shann Fuh. Complementary watermarks hiding for robust protection of images using DCT. In *Proceedings of International Symposium on Signal Processing and Intelligent System, ISSPIS '99*, Guangzhou City, China, November 1999.

[146] Chun-Shien Lu, Shih-Kun Huang, Chwen-Jye Sze, and Hong-Yuan Mark Liao. *A New Watermarking Technique for Multimedia Protection.* CRC Press, 2000.

[147] Chun-Shien Lu and Hong-Yuan Mark Liao. Oblivious watermarking using generalized gaussian. In *Proceedings of the 7th International Conference on Fuzzy Theory and Technology*, pages 260 – 263, Atlantic City, NJ, USA, February 2000.

[148] Chun-Shien Lu, Hong-Yuan Mark Liao, and Liang-Hua Chen. Multipurpose audio watermarking. In *Proceedings of the 15th International Conference on Pattern Recognition*, Barcelona, Spain, September 2000.

[149] Chun-Shien Lu, Hong-Yuan Mark Liao, Shih-Kun Huang, and Chwen-Jye Sze. Cocktail watermarking on images. In *Proceedings of the 3rd Information Hiding Workshop '99*, volume 1768, pages 333 – 347, Dresden, Germany, October 1999. Springer.

[150] Chun-Shien Lu, Hong-Yuan Mark Liao, Shih-Kun Huang, and Chwen-Jye Sze. Highly robust image watermarking using complementary modulations. In *Proceedings of the 2nd International Information Security Workshop*, pages 136 – 153, Kuala Lumpur, Malaysia, November 1999. Springer Verlag.

[151] Alessandra Lumini and D. Maio. A wavelet-based image watermarking scheme. In *Proceedings of the IEEE International Conference on Information Technology: Coding and Computing*, pages 122 – 127, Las Vegas, NV, USA, March 2000.

[152] Anne Manoury, Jacques Levy-Vehel, and Marie-Francoise Lucas. Watermarking d'images par paquets d'ondelettes. In *Proceedings of GRETSI '99*, Vannes, France, September 1999.

[153] Ross Martin and Douglas Cochran. Generalized wavelet transforms and the cortex transform. In *Proceedings of the 28th Asilomar Conference on Signals, Systems, and Computers*, November 1994.

[154] Lisa M. Marvel and Charles G. Boncelet. Capacity of the additive steganographic channel. *submitted to IEEE Transactions on Signal Processing*, June 1999.

[155] S. Maslakovic, I. R. Linscott, M. Oslick, and J. D. Twicken. Smooth orthonormal wavelet libraries: design and application. In *Proceedings of IEEE ICASSP '98*, pages 1793 – 1796, Seattle, WA, USA, May 1998.

[156] Kineo Matsui, Junji Ohnishi, and Yasuhiro Nakamura. Use of the wavelet transformation to embed signatures in images. *Systems and Computers in Japen*, 28(1):87 – 94, January 1997.

[157] Nicholas F. Maxemchuk and Steven Low. Marking text documents. In *Proceedings of the IEEE International Conference on Image Processing, ICIP '97*, volume 3, page 13, Santa Barbara, California, USA, October 1997.

[158] Nasir Memon and Jiri Fridrich. Further attacks on the yeung-mintzer fragile watermark. In Ping Wah Wong, editor, *Proceedings of IS&T/SPIE's 12th Annual Symposium, Electronic Imaging 2000: Security and Watermarking of Multimedia Content II*, volume 3971, San Jose, CA, USA, January 2000.

[159] Ruggero Milanese, Michel Cherbuliez, and Thierry Pun. Invariant content-based image retrieval using the fourier-mellin transform. In S. Singh, editor, *International Conference on Advances in Pattern Recognition, ICAPR '98*, pages 73 – 82, Plymouth, UK, November 1998. Springer Verlag.

[160] Matthew L. Miller and Jeffrey A. Bloom. Computing the probability of false watermark detection. In *Proceedings of the 3rd Information Hiding Workshop '99*, volume 1768, pages 146 – 158, Dresden, Germany, October 1999. Springer.

[161] Matthew L. Miller, Ingemar J. Cox, and Jeffrey A. Bloom. Watermarking in the real world: An application to DVD. In *Multimedia and Security Workshop at the 6th ACM International Multimedia Conference*, pages 71 – 76, Bristol, England, 1998.

[162] Frederick C. Mintzer, L. E. Boyle, A. N. Cazes, B. S. Christian, S. C. Cox, F. P. Giordano, Henry M. Gladney, Jong Chan Lee, M. L. Kelmanson, A. C. Lirani, Karen A. Magerlein, A. M. B. Pavani, and Fabio Schiattarella. Toward on-line, worldwide access to vatican library materials. *IBM Journal of Research & Development*, 40(2), 1995.

[163] Frederick C. Mintzer, Gordon W. Braudaway, and Minerva M. Yeung. Effective and ineffective digital watermarks. In *Proceedings of the IEEE International Conference on Image Processing, ICIP '97*, volume 3, page 9, Santa Barbara, California, USA, October 1997.

[164] Moni Naor and Adi Shamir. Visual cryptography. In Alfredo De Santis, editor, *Advances in Cryptology - EUROCRYPT '94, Workshop on the Theory and Application of Cryptographic Techniques*, volume 950 of *Lecture Notes in Computer Science*, pages 1 – 12, Perugia, Italy, 1994. Springer.

[165] Christian Neubauer and J. Herre. Digital watermarking and its influence on audio quality. In *Proceedings of the 105th Convention of the Audio Engineering Society*, September 1998.

[166] Gianluca Nicchiotti and E. Ottaviano. Non-invertible statistical wavelet watermarking. In *Proceedings of the 9th European Signal Processing Conference, EUSIPCO '98*, pages 2289 – 2292, Island of Rhodes, Greece, September 1998.

[167] Nikos Nikolaidis and Ioannis Pitas. Copyright protection of images using robust digital signatures. In *Proceedings of the IEEE International Conference on Accoustics, Speech and Signal Processing, ICASSP '96*, volume 4, pages 2168 – 2171, Atlanta, USA, May 1996.

[168] Nikos Nikolaidis and Ioannis Pitas. Digital image watermarking: An overview. In *Proceedings of the IEEE International Conference on Multimedia & Computing Systems, ICMCS '99*, volume 1, pages 1 – 6, Florence, Italy, June 1999.

[169] Ryutarou Ohbuchi, Hiroshi Masuda, and Masaki Aono. Watermarking three-dimensional polygonal models through geometric and topological modifications. In *IEEE Journal on Special Areas in Communications*, volume 16, pages 551 – 560, May 1998.

[170] Junji Ohnishi and Kineo Matsui. Embedding a seal into a picture under orthogonal wavelet transform. In *Proceedings of the IEEE Conference on Multimedia Computing and Systems*, pages 514 – 521, Hiroshima, Japan, June 1996.

[171] Junji Ohnishi and Kineo Matsui. A method of watermarking with multiresolution analysis and pseudo noise sequences. *Systems and Computers in Japan*, 29(5):11 – 19, May 1998.

[172] Job Oostveen, Ton Kalker, and Jean-Paul Linnartz. Optimal detection of multiplicative watermarks. In *Proceedings of the 10th European Signal Processing Conference, EUSIPCO '00*, Tampere, Finland, September 2000.

[173] Joseph J. K. O'Ruanaidh, F. M. Boland, and O. Sinnen. Watermarking digital images for copyright protection. In *Electronic Imaging and the Visual Arts, EVA '96*, Florence, Italy, 1996.

[174] Joseph J. K. O'Ruanaidh, W. J. Dowling, and F. M. Boland. Watermarking digital images for copyright protection. In *IEEE Conference on Vision, Image and Signal Processing, August 1996*, volume 143, pages 250 – 256, 1995.

[175] Joseph J. K. O'Ruanaidh and Thierry Pun. Rotation, scale and translation invariant spread spectrum digital image watermarking. *Signal Processing*, 66(3):303 – 317, May 1998.

[176] W. B. Pennebaker and J. L. Mitchell. *JPEG Still Image Compression Standard*. Van Nostrand Reinhold, New York, 1993.

[177] Shelby Pereira, Sviatoslav Voloshynovskiy, and Thierry Pun. Optimized wavelet domain watermark embedding strategy using linear programming. In Harold H. Szu, editor, *SPIE AeroSense 2000: Wavelet Applications VII*, Orlando, FL, USA, April 2000.

[178] Fabien A. P. Petitcolas and Ross J. Anderson. Evaluation of copyright marking systems. In *Proceedings of IEEE International Conference on Multimedia Computing and Systems '99*, volume 1, pages 574 – 579, Florence, Italy, June 1999.

[179] Fabien A. P. Petitcolas, Ross J. Anderson, and Markus G. Kuhn. Attacks on copyright marking systems. In David Aucsmith, editor, *Information hiding: second international workshop*, volume 1525 of *Lecture notes in computer science*, Portland, OR, USA, April 1998. Springer Verlag, Berlin, Germany.

[180] Birgit Pfitzmann. Information hiding terminology - results of an informal plenary meeting and additional proposals. In Ross Anderson, editor, *Information hiding: first international workshop*, volume 1174 of *Lecture Notes in Computer Science*, pages 347 – 350, Cambridge, UK, May 1996. Springer Verlag, Berlin, Germany.

[181] Birgit Pfitzmann and Michael Waidner. Asymmetric fingerprinting for larger collusions. In *4th ACM Conference on Computer and Communications Security*, Zurich, Switzerland, April 1997.

[182] Ioannis Pitas. A method for signature casting on digital images. In *Proceedings of the IEEE International Conference on Image Processing, ICIP '96*, volume 3, pages 215 – 218, Lausanne, Switzerland, September 1996. IEEE Press.

[183] Ioannis Pitas and T. H. Kaskalis. Applying signatures on digital images. In *Proceedings of the IEEE International Workshop on Nonlinear Signal and Image Processing*, pages 460 – 463, Marmaras, Greece, June 1995.

[184] Alessandro Piva, Mauro Barni, Franco Bartolini, and Vito Cappellini. DCT-based watermark recovering without resorting to the uncorrupted original image. In *Proceedings of the IEEE International Conference on Image Processing, ICIP '97*, volume 1, page 520, Santa Barbara, California, USA, October 1997.

[185] Alessandro Piva, Mauro Barni, Franco Bartolini, and Vito Cappellini. A watermarking technique for the protection of digital images IPR. In *Proceedings of EMMSEC, European Multimedia, Microprocessor System and Electronic Commerce Conference and Exhibition: Advances in Information Technologies: The Business Challenge*, pages 636 – 643, Florence, Italy, November 1997.

[186] Alessandro Piva, Mauro Barni, Franco Bartolini, and Vito Cappellini. Application-driven requirements for digital watermarking technology. In J. Y. Roger, editor, *Proceedings of EMMSEC, European Multimedia, Microprocessor System and Electronic Commerce Conference and Exhibition, Technologies for the Information Society: developments and Opportunities*, pages 513 – 520, Bordeaux, France, September 1998. IOS Press.

[187] Alessandro Piva, Roberto Caldelli, and Alessia DeRosa. A DWT-based object watermarking system for MPEG-4 video streams. In *Proceedings of the IEEE International Conference on Image Processing, ICIP '00*, Vancouver, Canada, September 2000.

[188] Christine I. Podilchuk. On tradsoffs in imperceptibility, robustness, and capacity. In *DFG V3D2 Watermarking Workshop 1999*, Erlangen, Germany, October 1999.

[189] Christine I. Podilchuk and Wenjun Zeng. Digital image watermarking using visual models. In Bernice E. Rogowitz, editor, *Proceedings of the 2nd SPIE Human Vision and Electronic Imaging Conference*, volume 3016, pages 100 – 111, June 1997.

[190] Christine I. Podilchuk and Wenjun Zeng. Image-adaptive watermarking using visual models. *IEEE Journal on Selected Areas in Communications, special issue on Copyright and Privacy Protection*, 16(4):525 – 539, May 1998.

[191] David Pollen. Parametrization of compactly supported wavelets. Technical report, Aware Inc., USA, 1989.

[192] Paolo Prandoni and Martin Vetterli. Perceptually hidden data transmission over audio signals. In *Proceedings of IEEE ICASSP '98*, Seattle, WA, USA, 1998.

[193] Joan Puate and Frédéric Jordan. Using fractal compression scheme to embed a digital signature into an image. In *Proceedings of the SPIE Visual Communications and Image Processing Conference*, volume 2915, pages 108 – 117, November 1996.

[194] Lintian Qiao and Klara Nahrstedt. Watermarking methods for MPEG encoded video: Towards resolving rightful ownership. In *Proceedings of the IEEE International Conference on Multimedia Computing and Systems, ICMCS '98*, pages 276 – 285, Austin, TX, USA, 1998.

[195] Mahalingam Ramkumar and Ali N. Akansu. A robust scheme for oblivious detection of watermarks / data hiding in still images. In *SPIE Symposium on Voice, Video and Data Communication*, pages 474 – 481, Boston, MA, USA, September 1998.

[196] Mahalingam Ramkumar and Ali N. Akansu. Capacity estimates for data hiding in compressed images. *IEEE Transactions on Image Processing*, 1999.

[197] Mahalingam Ramkumar and Ali N. Akansu. Self-noise suppression schemes for blind image steganography. In *Proceedings of SPIE: Multimedia Systems and Applications II*, volume 3845, Boston, MA, USA, September 1999.

[198] Mahalingam Ramkumar, Ali N. Akansu, and A. Aydin Alatan. On the choice of transforms for data hiding in compressed video. In *Proceedings of ICASSP '99*, volume 6, pages 3049 – 3052, Phoenix, AZ, USA, March 1999.

[199] Mahalingam Ramkumar, Ali N. Akansu, and A. Aydin Alatan. A robust data hiding scheme for images using DFT. In *Proceedings of the 6th IEEE International Conference on Image Processing, ICIP '99*, pages 211 – 215, Kobe, Japan, October 1999.

[200] K. R. Rao and Ping Yip. *Discrete cosine transform: algorithms, advantages, applications.* Academic Print, Boston, USA, 1990.

[201] Howard L. Resnikoff, Jun Tian, and Raymond O. Wells. Biorthogonal wavelet space: parametrization and factorization. *SIAM Journal on Mathematical Analysis*, August 1999.

[202] Amir Said and William A. Pearlman. A new, fast, and efficient image codec based on set partitioning in hierarchical trees. In *IEEE Transactions on Circuits and Systems for Video Technology*, volume 6, pages 243 – 250, June 1996.

[203] Yasuyuki Sakai, Hirokazu Ishizuka, and Kouichi Sakurai. A security of a watermarking for copyright protection using wavelet transform. *Transactions of Information Processing Society of Japan*, 38(12), December 1997.

[204] Diego Santa-Cruz and Touradj Ebrahimi. An analytical study of JPEG 2000 functionalities. In *Proceedings of the IEEE International Conference on Image Processing, ICIP '00*, Vancouver, Canada, September 2000.

[205] Diego Santa-Cruz and Touradj Ebrahimi. A study of JPEG 2000 still image coding versus other standards. In *Proceedings of the 10th European Signal Processing Conference, EUSIPCO '00*, Tampere, Finland, September 2000.

[206] Josef Scharinger. Robust watermark generation for multimedia copyright protection. In Markus Vincze, editor, *Robust Vision for Industrial Applications 1999*, pages 127 – 136. OCG, 1999.

[207] Josef Scharinger. Robust watermark generation for multimedia copyright protection. In *Proceedings of IWSSIP '99*, pages 177 – 180, Bratislava, Slovakia, 1999.

[208] J. Schneid and S. Pittner. On the parametrization of the coefficients of dilation equations for compactly supported wavelets. *Computing*, 51:165 – 173, May 1993.

[209] Gaspari Schotti. *Steganographica*. 1665.

[210] L. Schuchman. Dither signals and their effect on quantization noise. *IEEE Transaction on Communication Technology*, 12:162 – 165, December 1964.

[211] Sergio D. Servetto, Christine I. Podilchuk, and Kannan Ramchandran. Capacity issues in digital image watermarking. In *Proceedings of the IEEE International Conference on Image Processing, ICIP '98*, Chicago, IL, USA, October 1998.

[212] Jerome M. Shapiro. Embedded image coding using zerotrees of wavelet coefficients. *IEEE Transactions on Signal Processing*, 41(12):3445 – 3462, December 1993.

[213] Vassilios Solachidis, Nikos Nikolaidis, and Ioannis Pitas. Watermarking polygonal lines using fourier descriptors. In *Proceedings of the IEEE ICASSP 2000*, Istanbul, Turkey, June 2000.

[214] Vassilios Solachidis and Ioannis Pitas. Circularly symmetric watermark embedding in 2-d DFT domain. In *Proceedings of ICASSP '99*, volume 6, pages 3469 – 3472, Phoenix, AZ, USA, March 1999.

[215] Doug Stinson. Visual cryptography & threshold schemes. *Dr. Dobb's Journal*, 284:36 – 43, April 1998.

[216] Harold Stone. Analysis of attacks on image watermarks with randomized coefficients. Technical report, NEC Research Institute, May 1996.

[217] Jonathan K. Su, Joachim J. Eggers, and Bernd Girod. Capacity of digital watermarks subjected to an optimal collusion attack. In *Proceedings of the 10th European Signal Processing Conference, EUSIPCO '00*, Tampere, Finland, September 2000.

[218] Jonathan K. Su, Joachim J. Eggers, and Bernd Girod. Optimum attack on digital watermarks and its defense. In *34rd ASILOMAR Conference on Signals, Systems and Computers*, Asilomar, CA, USA, October 2000.

[219] Jonathan K. Su and Bernd Girod. Power-spectrum condition for energy-efficient watermarking. In *Proceedings of the 6th IEEE International Conference on Image Processing, ICIP '99*, Kobe, Japan, October 1999.

[220] Jonathan K. Su, Frank Hartung, and Bernd Girod. Digital watermarking of text, image and video documents. *Computers & Graphics*, 22(6):687 – 695, December 1998.

[221] Jonathan K. Su, Frank Hartung, and Bernd Girod. A channel model for a watermark attack. In Ping Wah Wong, editor, *Proceedings of the 11th SPIE Annual Symposium, Electronic Imaging '99, Security and Watermarking of Multimedia Contents*, volume 3657, San Jose, CA, USA, January 1999.

[222] Po-Chyi Su, Houng-Jyh Wang, and C.-C. Jay Kuo. Blind digital watermarking for cartoon and map images. In Ping Wah Wong, editor, *Proceedings of the 11th SPIE Annual Symposium, Electronic Imaging '99, Security and Watermarking of Multimedia Contents*, volume 3657, San Jose, CA, USA, January 1999.

[223] Po-Chyi Su, Houng-Jyh Wang, and C.-C. Jay Kuo. Digital image watermarking in regions of interest. In *The IS&T Image Processing, Image Quality, Image Capture, Systems Conference, PICS '99*, pages 295 – 300, Savannah, GA, USA, April 1999.

[224] Po-Chyi Su, Houng-Jyh Wang, and C.-C. Jay Kuo. Digital watermarking on EBCOT compressed images. In *Proceedings of SPIE's 44th Annual Meeting: Applications of Digital Image Processing XXII*, volume 3808, Denver, CO, USA, July 1999.

[225] Mitchell D. Swanson, Mei Kobayashi, and Ahmed H. Tewfik. Robust audio watermarking using perceptual masking. *Proceedings of the IEEE*, 86(6):1064 – 1087, June 1998.

[226] Mitchell D. Swanson and Ahmed H. Tewfik. A binary wavelet decomposition of binary images. *IEEE Transactions on Image Processing*, 5(12), December 1996.

[227] Mitchell D. Swanson, Bin Zhu, and Ahmed H. Tewfik. Multiresolution video watermarking using perceptual models and scene segmentation. In *Proceedings of the IEEE International Conference on Image Processing, ICIP '97*, volume 2, page 558, Santa Barbara, California, USA, October 1997.

[228] Mitchell D. Swanson, Bin Zhu, and Ahmed H. Tewfik. Multiresolution scene-based video watermarking using perceptual models. *IEEE Journal on Selected Areas in Communications*, 16(4):540 – 550, April 1998.

[229] Mitchell D. Swanson, Bin Zhu, and Ahmed H. Tewfik. Current state of the art - challenges and future directions for audio watermarking. In *Proceedings of the IEEE International Conference on Multimedia & Computing Systems, ICMCS '99*, Florence, Italy, June 1999.

[230] Wim Sweldens. The lifting scheme: A construction of second generation wavelets. *SIAM Journal of Mathemetical Analysis*, 29(2):511 – 546, March 1998.

[231] Weili Tang and Yoshinao Aoki. A haar wavelet transform in watermarking. In *The 1997 Society Conference of IEICE*, Tokyo, Japan, 1997.

[232] David Taubman. High performance scalable image compression with EBCOT. *IEEE Transactions on Image Processing*, 9(7):1158 – 1170, July 2000.

[233] Ahmed H. Tewfik, Mitchell D. Swanson, and Bin Zhu. Data embedding in audio: Where do we stand. In *Proceedings of ICASSP '99*, Phoenix, AZ, USA, 1999.

[234] Andrew Z. Tirkel and Charles F. Osborne. Image and watermark registration for monochrome and coloured images. In *Digital Image Computing, Technology and Applications*, pages 59 – 64, Wellington, New Zealand, 1997.

[235] Andrew Z. Tirkel, Charles F. Osborne, and Thomas E. Hall. Image and watermark registration. *Signal Processing*, (66):373 – 383, 1998.

[236] Johannes Trithemius. *Steganographia.* 1500.

[237] Min-Jen Tsai, Kuang-Yoo Yu, and Yi-Zhang Chen. Joint wavelet and spatial transformation for digital watermarking. *IEEE Transactions on Consumer Electronics*, 46:241 – 245, 1999.

[238] Min-Jen Tsai, Kuang-Yoo Yu, and Yi-Zhang Chen. Wavelet packet and adaptive spatial transformation of watemark for digital image authentication. In *Proceedings of the IEEE International Conference on Image Processing, ICIP '00*, Vancouver, Canada, September 2000.

[239] Sofia Tsekeridou and Ioannis Pitas. Embedding self-similar watermarks in the wavelet domain. In *Proceedings of the IEEE ICASSP 2000*, Istanbul, Turkey, June 2000.

[240] D. Tzovaras, N. Karagiannis, and M. G. Strintzis. Robust image watermarking in the subband or discrete cosine transform. In *Proceedings of the 9th European Signal Processing Conference, EUSIPCO '98*, pages 2285 – 2288, Island of Rhodes, Greece, September 1998.

[241] Pierre Vandergheynst, Martin Kutter, and Stefan Winkler. Wavelet-based contrast computation and application to digital image watermarking. In *Proceedings of the SPIE Wavelet Applications in Signal and Image Processing*, volume 4119, San Diego, CA, USA, July 2000.

[242] S. Voloshynovskiy, Alexander Herrigel, and Y. B. Rytsar. Watermark template attack. In *Proceedings of SPIE, Security and Watermarking of Multimedia Contents*, volume 4314, San Jose, CA, USA, January 2001.

[243] Sviatoslav Voloshynovskiy, Alexander Herrigel, Frédéric Jordan, Nazanin Baumgärtner, and Thierry Pun. A noise removal attack for watermarked images. In *Multimedia and Security Workshop at the 7th ACM International Multimedia Conference*, Orlando, FL, USA, October 1999.

[244] Sviatoslav Voloshynovskiy, Shelby Pereira, Alexander Herrigel, Nazanin Baumgärtner, and Thierry Pun. Generalized watermark attack based on watermark estimation and perceptual remodulation. In Ping Wah Wong, editor, *Proceedings of IS&T/SPIE's 12th Annual Symposium, Electronic Imaging 2000: Security and Watermarking of Multimedia Content II*, volume 3971, San Jose, CA, USA, January 2000.

[245] Sviatoslav Voloshynovskiy, Shelby Pereira, Victor Iquise, and Thierry Pun. Attack modelling: Towards a second generation benchmark, 2001.

[246] Sviatoslav Voloshynovskiy, Shelby Pereira, and Thierry Pun. Watermark attacks. In *DFG V3D2 Watermarking Workshop 1999*, Erlangen, Germany, October 1999.

[247] George Voyatzis and Ioannis Pitas. Application of toral automorphisms in image watermarking. In *Proceedings of the IEEE International Conference on Image Processing, ICIP '96*, volume 2, pages 237 – 240, Lausanne, Switzerland, September 1996. IEEE Press.

[248] George Voyatzis and Ioannis Pitas. Digital image watermarking using mixing systems. *Computer & Graphics*, 22(4):405 – 416, August 1998.

[249] Steve Walton. Image authentication for a slippery new age. *Dr. Dobb's Journal*, (229):18 – 26, April 1995.

[250] Houng-Jyh Wang, Yi-Liang Bao, C.-C. Jay Kuo, and Homer Chen. Multi-threshold wavelet codec (MTWC). Technical report, Derpartment of Electrical Engineering, University of Southern California, Los Angeles, CA, USA, Geneva, Switzerland, March 1998.

[251] Houng-Jyh Wang and C.-C. Jay Kuo. High fidelity image compression with multithreshold wavelet coding (MTWC). In *SPIE's Annual meeting - Application of Digital Image Processing XX*, San Diego, CA, USA, August 1997.

[252] Houng-Jyh Wang and C.-C. Jay Kuo. Image protection via watermarking on perceptually significant wavelet coefficients. In *Proceedings of the IEEE Workshop on Multimedia Signal Processing*, Los Angeles, CA, USA, December 1998. IEEE.

[253] Houng-Jyh Wang and C.-C. Jay Kuo. An integrated approach to embedded image coding and watermarking. In *Proceedings of IEEE ICASSP '98*, Seattle, WA, USA, May 1998.

[254] Houng-Jyh Wang and C.-C. Jay Kuo. Watermark design for embedded wavelet image codec. In *Proceedings of the SPIE's 43rd Annual Meeting, Applications of Digital Image Processing*, volume 3460, pages 388 – 398, San Diego, CA, USA, July 1998.

[255] Houng-Jyh Wang, Po-Chyi Su, and C.-C. Jay Kuo. Wavelet-based blind watermark retrieval technique. In *Proceedings of the SPIE Phonics East Symposium on Voice, Video and Data Communications*, volume 3528, pages 440 – 451, Boston, MA, USA, November 1998.

[256] Houng-Jyh Wang, Po-Chyi Su, and C.-C. Jay Kuo. Wavelet-based digital image watermarking. *Optics Express*, 3(12):497, December 1998.

[257] Andrew B. Watson. The cortex transform: rapid computation of simulated neural images. *Computer Vision, Graphics, and Image Processing*, 39(3):311 – 327, 1987.

[258] Andrew B. Watson, Gloria Y. Yang, Joshua A. Solomon, and John Villasenor. Visual thresholds for wavelet quantization error. In B. Rogowitz, editor, *Proceedings of the SPIE*, volume 2657, pages 382 – 392, 1996.

[259] Andrew B. Watson, Gloria Y. Yang, Joshua A. Solomon, and John Villasenor. Visibility of wavelet quantization noise. *IEEE Transaction in Image Processing*, 6:1164 – 1175, 1997.

[260] Mladen Victor Wickerhauser. *Adapted Wavelet Analysis from Theory to Software.* A. K. Peters, Ltd., 1993.

[261] Raymond B. Wolfgang, Christine I. Podilchuk, and Edward J. Delp. The effect of matching watermark and compression transforms in compressed color images. In *Proceedings of the IEEE International Conference on Image Processing, ICIP '98*, Chicago, IL, USA, October 1998.

[262] Raymond B. Wolfgang, Christine I. Podilchuk, and Edward J. Delp. Perceptual watermarks for digital images and video. *Proceedings of the IEEE, Special Issue on Identification and Protection of Multimedia Information*, 87(7):1108 – 1126, July 1999.

[263] Chuan-Fu Wu and Wen-Shyong Hsieh. Multiresolution watermarking technique with wavelet based stack-run coding. In *Proceedings of the 7th International Conference on Fuzzy Theory and Technology*, page 374, Atlantic City, NJ, USA, February 2000.

[264] Min Wu and Bede Liu. Attacks on digital watermarks. In *33rd ASILOMAR Conference on Signals, Systems and Computers*, Pacific Grove, CA, USA, 1999.

[265] Xiang-Gen Xia, Charles G. Boncelet, and Gonzalo R. Arce. A multiresolution watermark for digital images. In *Proceedings of the IEEE International Conference on Image Processing, ICIP '97*, volume 1, page 548, Santa Barbara, California, USA, October 1997.

[266] Xiang-Gen Xia, Charles G. Boncelet, and Gonzalo R. Arce. Wavelet transform based watermark for digital images. *Optics Express*, 3(12):497, December 1998.

[267] Liehua Xie and Gonzalo R. Arce. Blind wavelet based digital signature for image authentication. In *Proceedings of the 9th European Signal Processing Conference, EUSIPCO '98*, pages 21 – 24, Island of Rhodes, Greece, September 1998.

[268] Liehua Xie and Gonzalo R. Arce. Joint wavelet compression and authentication watermarking. In *Proceedings of the IEEE International Conference on Image Processing, ICIP '98*, Chicago, IL, USA, 1998.

[269] Wenjun Zeng, Scott Daly, and Shawmin Lei. Point-wise extended visual masking for JPEG2000 image compression. In *Proceedings of the IEEE International Conference on Image Processing, ICIP '00*, Vancouver, Canada, September 2000.

[270] Wenjun Zeng, Scott Daly, and Shawmin Lei. Visual optimization tools in JPEG2000. In *Proceedings of the IEEE International Conference on Image Processing, ICIP '00*, Vancouver, Canada, September 2000.

[271] Wenjun Zeng and Shawmin Lei. Transform domain perceptual watermarking with scalable visual detection- a proposal for JPEG2000. Technical report, Digital Video Department, Sharp Laboratories of America, Inc., Camas, WA, USA, March 1998.

[272] Jian Zhao. Copyright protection technologies in the digital world. Technical report, Fraunhofer Institute for Computer Graphics, Darmstadt, Germany, 1996.

[273] Jian Zhao. A WWW service to embed and prove digital copyright watermarks. In *European Conference on Multimedia Applications, Services and Techniques, ECMAST '96*, volume 2, pages 695 – 710, Louvain-la-Neuve, Belgium, May 1996.

[274] Jian Zhao and Eckhard Koch. Embedding robust labels into images for copyright protection. In *Proceedings of the International Congress on Intellectual Property Rights for Specialized Information, Knowledge and New Technologies*, pages 242 – 251, Vienna, Austria, August 1995.

[275] Wenwu Zhu, Zixiang Xiong, and Ya-Qin Zhang. Multiresolution watermarking for images and video: a unified approach. In *Proceedings of the IEEE International Conference on Image Processing, ICIP '98*, pages 465 – 468, Chicago, IL, USA, October 1998.

[276] Wenwu Zhu, Zixiang Xiong, and Ya-Qin Zhang. Multiresolution watermarking for images and video. *IEEE Transactions on Cirtuits and Systems for Video Technology*, 9(4):545 – 550, June 1999.

[277] H. Zou and Ahmed H. Tewfik. Parametrization of compactly supported orthonormal wavelets. *IEEE Transactions on Signal Processing*, 41(3):1423 – 1431, March 1993.

Appendix A - JPEG2000

The upcoming next-generation image coding standard, JPEG2000, will be based on the DWT. Many new requirements and desirable features [27] that were taken into consideration for the new compression standard, such as

- coding performance, especially at high-compression rates (beyond a ratio of 1 : 32),
- progressive transmission,
- region-of-interest (ROI) coding capabilities,
- resolution scalability and perceptual quality scalability and
- reliability and security.

As of this writing (December 2000), the final committee draft (FCD) of the ISO document 15444 has been made available. In [204], the new image coding standard is compared to previous standards such as JPEG, JPEG-LS (lossless coding) and progressive JPEG.

Security and watermarking has been an issue during the JPEG2000 standardization process[3]. Several proposals have been made to incorporate watermark embedding into the new coding standard, for example Wang [253] and Zeng [271]. The JPEG2000 file format ("`.j2k`") provides an optional header tag that can be used to plug-in watermarking technology.

Although JPEG2000 is based on the wavelet transform, it is quiet different from other prominent wavelet image compression techniques such as EZW [212] or SPIHT [202] since JPEG2000 operates on independent code-blocks, much like the EBCOT [232] algorithm. First watermarking schemes that take into account block-based coding in the wavelet domain have been proposed by Su [224].

For a brief discussion of the mentioned image compression algorithms can be found in section 2.7.3.

[3] See the documents and resentation of the seminar on image security, `http://eurostill.epfl.ch/~ebrahimi/JPEG2000.htm`, Vancouver, Canada, July 1999.

Appendix B - Test Images

The following test images were selected and are shown on the following pages. All images are 512 by 512 pixels large, 8 bits per pixel (bpp) gray-scale.

Lena The digitized Playboy center-fold[4], Miss November 1972, is the classical image in image processing. The smooth regions of her shoulder and the sharp contrast with the background makes it difficult to embed a watermark without adding visible distortion. On the other hand, it is easy to manipulate the textured area of the feather.

Baboon The Baboon image contains a lot of texture and is therefore relatively easy to watermark.

Goldhill The Goldhill image has a lot of small detail and is therefore an easy host image. The smooth background can be tricky, however.

Fishing Boat The Fishing Boat image is a good target to tamper with. The lighthouse, the name of the boat and the man in the foreground can be easily removed, compare with section 4.4.2.2.

Cameraman The large, smooth background of the Cameraman image makes it very difficult to achieve watermark capacity and imperceptibility.

[4] See `http://www.image.cityu.edu.hk/images/lenna/Lenna97.html` for a more complete version of the Lena (or Lenna) story.

Figure 21: Lena, 512×512 gray-scale image, 8 bpp.

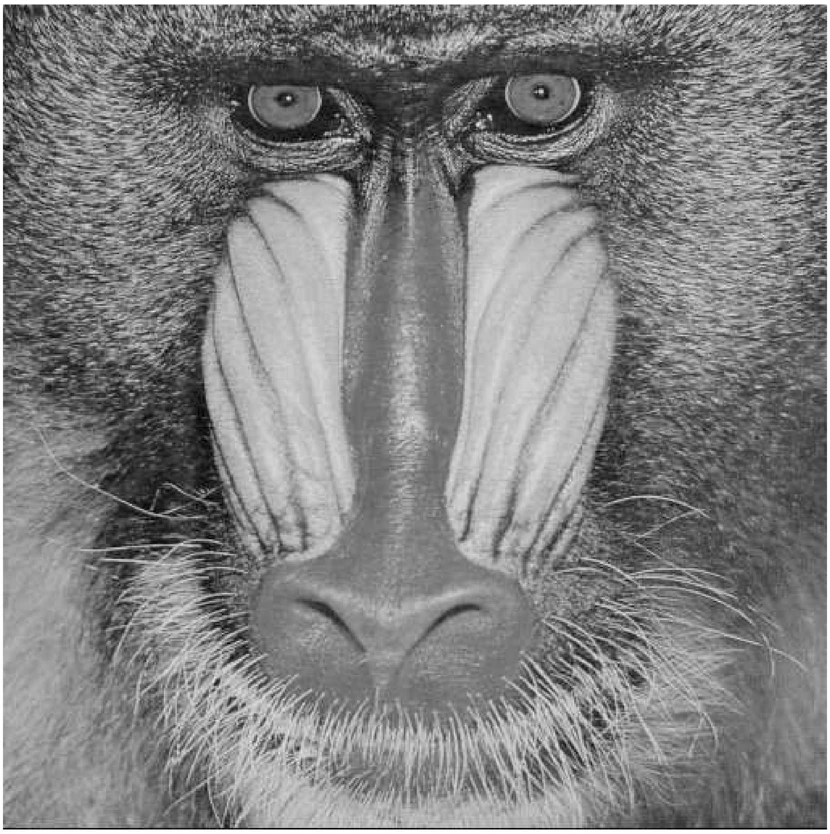

Figure 22: Baboon, 512×512 gray-scale image, 8 bpp.

Figure 23: Goldhill, 512×512 gray-scale image, 8 bpp.

Figure 24: Fishing Boat, 512×512 gray-scale image, 8 bpp.

Figure 25: Cameraman, 512×512 gray-scale image, 8 bpp.

Curriculum Vitae

Name Peter Meerwald

Date of birth April 18, 1975

Place of birth Salzburg, Austria

Parents Ingrid & Franz Meerwald

Education

1981 - 1985	Volksschule Aigen (primary school), Salzburg
1985 - 1989	Bundesrealgymnasium Akademiestraße (secondary school), Salzburg
1989 - 1994	Bundeshandelsakademie II (commercial high school), Salzburg
1995	Zivildienst (community service), Lebenshilfe Salzburg

Studies

1994 - 1998	Computer Science, University of Salzburg Undergraduate Program
1995 -	Legal Studies, University of Salzburg
1998 - 1999	Computer Science, Bowling Green State University, Ohio, USA Graduate Program (Master of Science, August 1999)
1998 -	Computer Science, University of Salzburg Graduate Program