Attacks, applications and evaluation of known watermarking algorithms with Checkmark

Peter Meerwald[a], Shelby Pereira[b]
[a]Dept. of Scientific Computing, University of Salzburg, Austria
[b]CUI, University of Geneva, Switzerland

ABSTRACT

The Checkmark benchmarking tool was introduced to provide a framework for application-oriented evalutaion of watermarking schemes. In this article we introduce new attacks and applications into the existing Checkmark framework.

In addition to describing new attacks and applications, we also compare the performance of some well-known watermarking algorithms (proposed by Bruyndonckx,[1,2] Cox,[3] Fridrich,[4] Dugad,[5] Kim,[6] Wang,[7] Xia,[8] Xie,[9] Zhu[10] and Pereira[11–13]) with respect to the Checkmark benchmark. In particular, we consider the 'non-geometric' application which contains tests that do not change the geometry of image. This attack constraint is artificial, but yet important for research purposes since a number of algorithms may be interesting, but would score poorly with respect to specific applications simply because geometric compensation has not been incorporated. We note, however, that with the help of image registration,[14] even research algorithms that do not have counter-measures against geometric distortion[15] – such as a template or reference watermark – can be evaluated.

In the first version of the Checkmark benchmarking program, application-oriented evaluation was introduced, along with many new attacks not already considered in the literature. A second goal of this paper is to introduce new attacks and new applications into the Checkmark framework. In particular, we introduce the following new applications: video frame watermarking, medical imaging and watermarking of logos. Video frame watermarking includes low compression attacks and distortions which warp the edges of the video as well as general projective transformations which may result from someone filming the screen at a cinema. With respect to medical imaging, only small distortions are considered and furthermore it is essential that no distortions are present at embedding. Finally for logos, we consider images of small sizes and particularly compression, scaling, aspect ratio and other small distortions. The challenge of watermarking logos is essentially that of watermarking a small and typically simple image. With respect to new attacks, we consider: subsampling followed by interpolation, dithering and thresholding which both yield a binary image.

Keywords: Watermarking, attacks, benchmarking, perceptual metric

1. INTRODUCTION

Digital watermarking has emerged as an appropriate tool for the protection of author's rights. It is now well accepted that an effective watermarking scheme must successfully deal with the triple requirement of *imperceptibility* (visibility) - *robustness* - *capacity*.[16] *Imperceptibility* requires that the marked data and the original data should be perceptually undistinguishible. *Robustness* refers to the fact that the embedded information should be reliably decodable after alterations of the marked data. Often the level of robustness is dictated by the application. *Capacity* refers to the amount of information that is being embedded in the watermark. In typical applications we require between 60 and 100 bits. This is necessary so as to uniquely associate images with buyers and sellers.

In addition to these requirements, the issue of algorithm complexity is also of importance. In some applications for example, it is necessary that the algorithms lend themselves to a hardware implementation. In other applications such as video watermarking, real-time embedding and detection may be essential. To further complicate the issue, the requirement on complexity may depend on the protocols used to distribute the media.

[a]Department of Scientific Computing, University of Salzburg, Jakob-Haringer-Str. 2, A-5020 Salzburg, Austria. pmeerw@cosy.sbg.ac.at; [b]CUI, University of Geneva, 24 Rue General Dufour, 1211 Geneva 4, Switzerland. shelby.pereira@cui.unige.ch

Given the relatively complex tradeoffs involved in designing a watermarking system, the question of how to perform fair comparisons between different algorithms naturally arises. A lack of systematic benchmarking of existing methods however creates confusion amongst content providers and watermarking technology suppliers. The benchmarking tool Stirmark[17,18] integrates a number of image processing operations or geometrical transformations aimed at removing watermarks from a stego image. The design of this tool does not take into account the statistical properties of the images and watermarks in the design of attacks. As a result, pirates can design more efficient attacks that are not currently included in the benchmarking tools. This could lead to a tremendous difference between what existing benchmarks test and real world attacks. Another problem with the Stirmark benchmarking tool is that it does not take into account the fact that different applications require different levels of robustness.

The Checkmark benchmarking tool[19] provides an application-oriented environment for the benchmarking of watermarking algorithms which overcomes many of the limitations of the Stirmark package. Furthermore, an important step was taken in designing attacks which include important priors on the image and the watermarking algorithm.[20] In particular, estimation based techniques were used to derive the optimal attacks for a given watermark distribution. Furthermore, a new method based on Watson's metric[21] was proposed for determining visual quality of a watermark image. It is the aim of this paper to extend Checkmark with new attacks and applications and to evaluate a set of well known algorithms against the benchmark.

The paper is structured as follows. In section 2 we review and categorize existing attacks. In section 3 we present the new attack types. In section 4 we describe new attacks recently implemented into Checkmark 1.2. We then turn our attention to new applications in section 5.

Finally in section 6 we present our results followed by a conclusion in section 7.

2. STATE-OF-ART WATERMARKING ATTACKS

We will adopt the attack classification scheme detailed in our previous paper[19] which we extend to include new attacks included in the current version (Checkmark 1.2). The wide class of existing attacks can be divided into four main groups: interference and removal attacks, geometrical attacks, cryptographic attacks and protocol attacks. Figure 1 summarizes the different attacks.

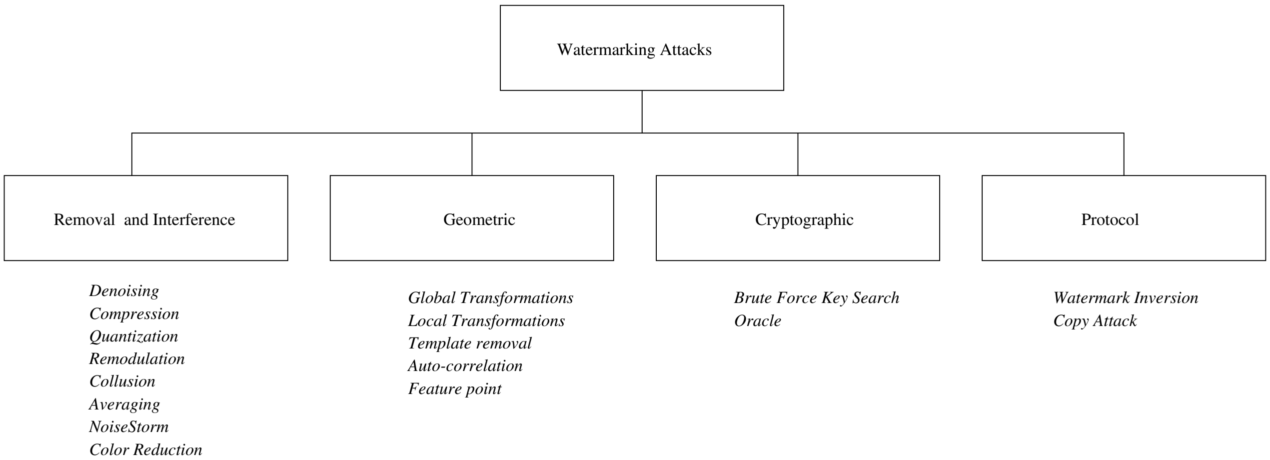

Figure 1. Classification of watermarking attacks

Organization of Checkmark Attacks. The Checkmark program uses a hierarchical organization of attacks which follows directly from figure 1. In particular, the Checkmark attack hierarchy contains four levels: attack groups, attack types, attack classes, and individual attacks. The attack group level contains the four broad categories: interference and removal, geometric, cryptographic, and protocol. Within each group we have attack types. For example in the removal and interference group one attack type is denoising. The attack types are further subdivided into classes. If we continue the example, the denoising attack type contains the ML and MAP attack classes. Finally, each class contains one or more individual attacks. For example the MAP attack class contains Wiener filtering, soft thresholding and hard thresholding.

3. NEW ATTACK TYPES

3.1. Synchronization Type Attacks

The synchronization type attacks are those which attack the synchronization method being used by the watermark detector. Three classes of attacks fall within this type: template removal, auto-correlation attack, and feature point attack. We consider these in order.

The template removal class now includes the well known attack on the DFT template.[20] The attack consists of removing peaks in the DFT domain and then slightly rotating the image. DFT peaks have been used to synchronize algorithms since they can be detected after linear transformations.[22]

The auto-correlation attack consists of: computing the auto-correlation function of the estimated watermark, detecting significant peaks in the auto-correlation, and finally using this information to attack the watermark. The peaks in the auto-correlation function result from periodicities in the watermark. If the period of the watermark can be determined, the attacker can correctly estimate the watermark by exploiting the redundancy in a straightforward way. Once the watermark has been estimated, a certain percentage (about 30 percent) of the values can be flipped and then added to the watermark image thereby severely weakening the original watermark.[20]

Feature points in images have also been used to create invariant watermarking schemes.[23,24] One possible attack is to modify the image in the locations of feature points. This is similar to performing a localized random bending of the image. For typical schemes, a shift of a few pixels is enough to desynchronize the algorithm and render the watermark un-detectable.

Checkmark 1.2 includes only the template removal attack. We will include the auto-correlation and feature point attack in a future release.

3.2. Color Reduction Attack Type

The color reduction attack type includes two major classes: color-to-grayscale reduction, grayscale-to-black and white. The color-to-gray scale class includes attacks from any color space to a grayscale space in which pixels values lie between 0 and 255. The grayscale-to-black and white class includes attacks such as dithering, thresholding and halftoning.

4. NEW ATTACKS

Having reviewed the attacks contained in our second generation benchmark, we now consider a number of new attacks which are contained in the current implementation. None of the attacks which we describe in this section have been included in the Stirmark benchmarking tool or in version 1.0 of the Checkmark program.

4.1. Downsampling and Interpolation

The downsampling followed by interpolation (upsampling) attack falls in the geometric attack group. In particular, we consider cases where the image is downsampled by a factor of two and then upsampled by a factor of two using bicubic interpolation. A good watermarking scheme should be able to recover the watermark despite the interpolation artifacts. We also include a more powerful attack where the image is downsampled by a factor of two and then upsampled by a factor of 1.9 or 2.1. The effect of this is that the attacked image is no longer of the same size as the original and consequently the watermark detector must be able to synchronize itself despite the interpolation artifacts. Other values of the downsampling and upsampling parameters are considered, however, we restrict ourselves to the case where downsampling is performed first since this necessarily leads to a greater loss of information. An example image along with the resampled image appear in figure 2. Although the quality of the attacked image remains good, in most cases the watermark is rendered undecodable.

4.2. Color Reduction Attacks

We include two new color reduction attacks namely dithering and thresholding. Dithering is important since it is a standard technique used to generate black-and-white images from grayscale images while maximizing perceived quality. We also consider thresholding although typically the resulting images are of poor quality. In order to choose the thresholding level, we first compute a histogram of the pixels and set the level to be the midpoint. We are currently investigating halftoning and inverse halftoning attacks which will be incorporated into future versions of the program. An example of the dithering and thresholding attacks appear in figure 3.

(a) original image (b) resampled image

Figure 2. Original image and image after downsampling/upsampling

(a) dithered image (b) thresholded image

Figure 3. Dithering and threshold attack

5. NEW APPLICATIONS

Checkmark 1.0 included the following applications: copyright protection, banknote protection and the artificial non-geometric application. In Checkmark 1.2, we have now added three new applications to Checkmark which we describe below.

Logo: We consider attacks on logos which are typically of small and arbitrary sizes. Attacks in this category are similar to those in the non-geometric category, however we are also interested in scaling and aspect ratio changes as well as a small cropping around the border. We have included a special set of logo images for this application. The images appear in figure 4 and contain small images which are typically not square.

Medical Images: Here we only consider copyright protection of medical images and *not* the authentication application which is equally important. In many instances, databases of medical images may require copyright protection. Relevant attacks in this application are limited to those which do not noticeably modify the image since any noticeable modification will render the image useless. Many medical images are x-ray type images which large amounts of black areas and then large concentrations of very bright pixels. Watermark embedders should be able to handle the unusual distribution of pixels. Sample images included in Checkmark 1.2 appear in figure 5.

Video: This application considers those attacks relevant to an individual *frame* of a video. In particular, projective transformations are interesting since they represent situations in which a person records a movie with a

Figure 4. Sample logo images used in Checkmark 1.2

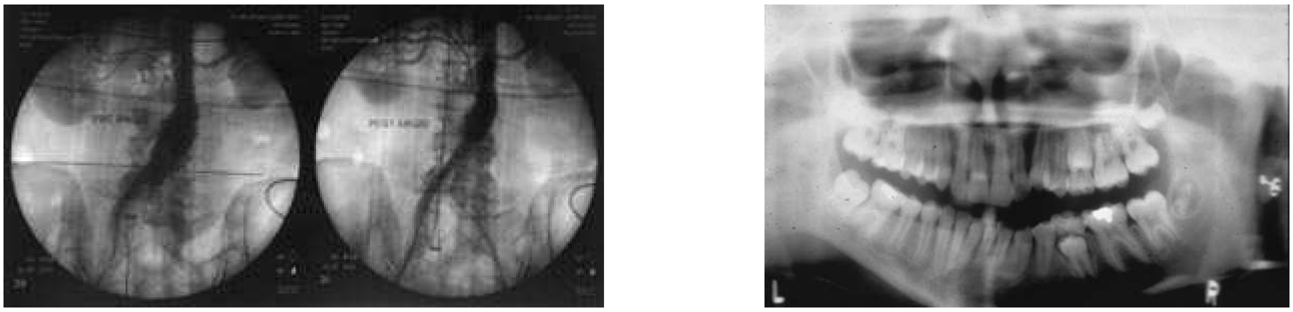

Figure 5. Sample medical images used in Checkmark 1.2

hand-held digital camera. Furthermore in video applications, large amounts of compression may be encountered in order to reduce the bandwidth requirements. Future versions of Checkmark will include actual attacks on MPEG streams.

6. RESULTS

Checkmark version 1.2 is available from `http://watermarking.unige.ch/Checkmark`. The program is written in Matlab and contains an XML description of the relationship of a given application to a set of attacks. Consequently, rather than proposing an overall score, it is easy to generate the results as a function of application. The use of XML yields a flexible way of adding new attacks and applications and weighting the results as a function of an application. Another advantage of XML is that the results can then be easily parsed and converted from one format to another.

Results for various algorithms are now being posted at `http://watermarking.unige.ch/Checkmark` and mirrored at the central watermarking site `www.watermarkingworld.org` along with references to the associated publications. Consequently results can be easily compared. In order to have your results included in the tables, the XML result files generated by Checkmark should be emailed to shelby.pereira@cui.unige.ch.

6.1. Benchmarked Algorithms

The algorithms we compare have been implemented closely following the descriptions in the corresponding papers. The source code of these implementations is available from `http://www.cosy.sbg.ac.at/~pmeerw/Watermarking`. The algorithms have been selected to capture a wide range of different embedding techniques (i.e. additive embedding vs. quantization-based strategies) and transform domains (i.e. spatial, DCT and DWT domain). Checkmark provides quality metrics (wPSNR and Watson) that allow to weight the detector performance according to the unobstrusiveness of the embedded mark. In particular, lower weightings are given to images which are more distorted.

This makes it possible to compare a wide range of watermarking algorithms on a fair basis. The results are divided into two groups: blind and non-blind. We will briefly characterize the groups and the present the main ideas of the algorithms belonging to these two groups.

Blind Schemes: Results for algorithms which do not require the original image at detection time. These are the most important since in practical situations we do not want to search for the original image in order to detect the watermark. In this group, we have investigated the following watermarking algorithms:

- **Bruyndonckx**[1,2] scheme is a spatial domain algorithm that embeds one bit of information in each selected image block. Two regions in a block are modified to impose a luminance difference between these regions. Block categorization and selection is used to increase robustness.
- **Dugad's**[5] algorithm modifies significant coefficients of the wavelet-domain detail subbands. Correlation detection is performed separately for each subband and decomposition level.
- Operating in the low- and mid-frequency range of the DCT domain, **Fridrich's**[4] algorithm aims to combine the advantages of robustness and imperceptability.
- An algorithm proposed by **Xie**[9] quantizes the coefficients of the approximation image in the wavelet-domain. To embed one bit of information, a coefficient triple of a non-overlapping 3×1 sliding window is selected and the median coefficient is manipulated.
- **Pereira's**[11–13] algorithms are more complex than the other schemes considered here. They involve a template mark to counter geometric attacks, employ error-correction codes and perform sophisticated masking.

Non-Blind Schemes: Results for algorithm which require the original image at detection time. In practice, a search in a large database of images may be required to find the image. Consequently in most situations blind schemes are preferred. The following algorithms were considered in this group:

- The **Cox**[3] algorithm adds a pseudo-random noise sequence to the 1000 largest DCT coefficients.
- **Kim's**[6] scheme modifies the coefficients in the approximation and detail subbands in the DWT domain. The energy of the watermark is adapted to the decomposition level.
- **Wang's**[7] algorithm selects significant coefficients in the perceptually most important subbands of an image in the wavelet domain and adds a pseudo-random noise sequence.
- The **Xia**[9] and **Zhu**[10] scheme both select significant wavelet coefficients from the detail subbands and add, again, a pseudo-random noise vector. The algorithms differ in the way the noise vector in generated and the implicit masking technique used.

Note that the payload embedded by the above algorithms ranges from one bit (e.g. Cox, Xia, Kim, Wang, Zhu) to more than hundred bits (e.g. Bruyndonckx and Xie), other schemes lie somewhere in between (e.g. Dugad, Fridrich, Pereira). To keep the algorithms simple and to make the results better comparable, we modified the detector of the algorithms such that the output is always a yes/no decision. The notable exception are the schemes by Pereira (i.e. UniGe1999 and UniGe2000) which perform reliable detection of 64 and 56 bits, respectively, making use of error-corrective coding.

In order to compute a yes/no decision, the Hamming distance between the extracted bit sequence and the embedded sequence was computed and compared with a threshold.

The problem of comparing watermarking schemes embedding a message versus a one-bit signature was further elaborated by Solachidis[25] and requires the evaluation of the Receiver Operating Characteristic (ROC). This leads to testing multiple messages and multiple keys for each image, watermarking algorithm and attack type. Hence, the computational requirements are dramatically increased which is out-of-scope for our present work.

(a) additive noise, wPSNR 35.2dB

(b) Cox algorithm, wPSNR 35.4dB

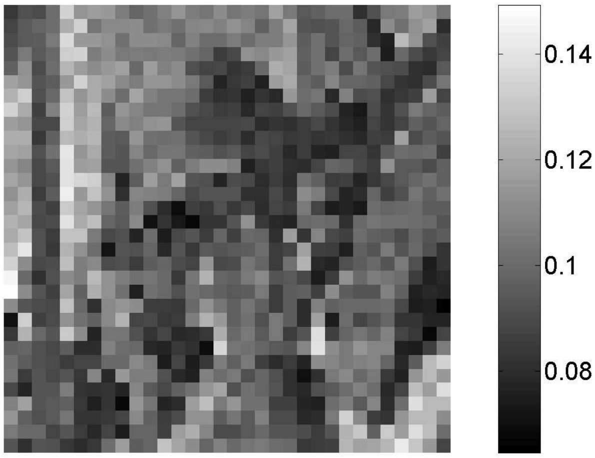

(c) Watson perceptual mask, TPE 0.098

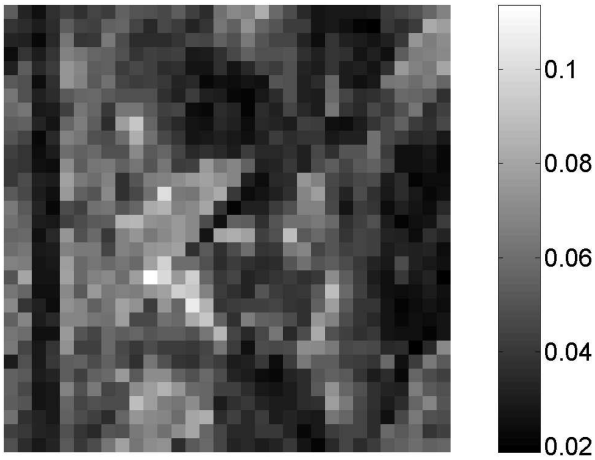

(d) Watson perceptual mask, TPE 0.047

Figure 6. Comparing additive noise and the Cox scheme at a PSNR of 34dB using the Watson perceptual metric

6.2. Quality Metrics

The Checkmark package includes PSNR, weighted PSNR (wPSNR) and the Watson metric which comprises the TPE, the total perceptual error, and the NLPE, the number of local blocks with visible artifacts due to watermark embedding. We illustrate the performance of the three metrics with the main goal of demonstrating the weaknesses of the commonly used PSNR metric.

Example 1 compares a non-adaptive embedding scheme with the Cox scheme[3] at 34dB PSNR to show the importance of masking and the inadequacy of PSNR. We compare embedding additive noise without adapting to the image with the Cox algorithm which modifies certain perceptual significant coefficients in the DCT domain (thereby performing implicit masking) – see figure 6. In both cases, the PSNR is set to 34db, however for the non-adaptive algorithm, the watermark is clearly visible. Watson's NLPE metric indicates that 91 blocks (16x16 coefficients) are visible. Also the total perceptual error (TPE) is 0.098 for the non-adaptive case while it is only 0.0469 for the Cox algorithm. This indicates that Watson's metric is indeed effective in this case for identifying a too visible watermark. Note that in both cases the wPSNR is about the same. Also note that the Watson metric correctly identifies the visible regions.

Algorithm		Image 1	Image 2	Image 3	Image 4	Image 5
Bruyndonckx	TPE	0.01	0.0	0.01	0.0	0.01
	NLPE	0	0	0	0	0
	wPSNR	48.87	54.79	53.87	56.99	55.07
Dugad	TPE	0.06	0.03	0.02	0.06	0.03
	NLPE	13	15	3	54	5
	wPSNR	42.07	42.84	44.36	44.62	44.37
Fridrich	TPE	0.11	0.09	0.09	0.09	0.08
	NLPE	75	151	188	140	106
	wPSNR	27.03	30.66	28.8	31.81	30.67
Xie	TPE	0.02	0.01	0.02	0.01	0.01
	NLPE	4	2	0	0	0
	wPSNR	33.85	41.87	36.069	42.88	43.97

Table 1. Perceptual quality measures on the images watermarked with the blind algorithms

Algorithm		Image 1	Image 2	Image 3	Image 4	Image 5
Cox	TPE	0.03	0.02	0.02	0.02	0.02
	NLPE	0	0	0	0	0
	wPSNR	40.65	41.73	38.32	46.45	43.78
Kim	TPE	0.01	0.01	0.00	0.01	0.01
	NLPE	0	0	0	0	0
	wPSNR	60.15	59.46	60.71	59.90	60.33
Wang	TPE	0.02	0.01	0.01	0.01	0.01
	NLPE	0	0	0	0	0
	wPSNR	47.47	47.80	46.19	51.69	49.01
Xia	TPE	0.05	0.03	0.02	0.06	0.02
	NLPE	15	7	2	73	5
	wPSNR	45.42	46.86	49.98	45.77	49.206
Zhu	TPE	0.03	0.02	0.02	0.02	0.02
	NLPE	0	2	0	0	0
	wPSNR	44.17	45.020	42.52	48.27	46.10

Table 2. Perceptual quality measures on the images watermarked with the non-blind algorithms

Example 2 compares a non-adaptive scheme with the Fridrich scheme[4] at a very low PSNR of 27dB to once again show the importance of masking and the inadequacy of PSNR. Both images appear distorted, however, for the non-adaptive algorithm, the watermark is clearly visible. Watson's metric indicates that 1024 blocks (16x16 coefficients) are visible (in other words, all blocks contain a visible watermark). Also the total perceptual error (TPE) is 0.2 for the non-adaptive case while it is only 0.105 for the Fridrich algorithm. In fact, even with the low PSNR, the Fridrich algorithm provides a reasonable quality image even though the watermark is visible in some parts.

In tables 1 and 2 we evaluate the perceptual quality of the watermarked images using the Watson metric for the blind and non-blind algorithms, respectively. Please note that not all schemes allow for easy control of the embedding strength. In these cases, the parameters were set to the closest values that give acceptable results (determined by looking at the images).

6.3. Attack Results

The overall attack results for the different Checkmark applications are shown in tables 3 and 4, respectively. The numbers quoted are percentages of correct detections. Five images were tested – the total number of attacks is indicated below each application name.

Algorithm (bits) (images, attacks)	Non-Geometric (5, 230)	Copyright (5, 1910)	Banknote (5, 890)	Logo (5, 385)	Medical (5, 225)	Video (5, 965)
Fridrich (1)	70%	N/A	N/A	N/A	N/A	N/A
Dugad (1)	88%	N/A	N/A	N/A	N/A	N/A
Bruyndonckx (1)	61%	N/A	N/A	N/A	N/A	N/A
Xie (1)	93%	N/A	N/A	N/A	N/A	N/A
UniGe1999 (64)	N/A	39%	47%	N/A	30%	N/A
UniGe2000 (56)	61%	N/A	N/A	13%	N/A	14%

Table 3. Checkmark 1.2 results for the blind algorithms; N/A indicates that data is unavailable, or the algorithm is not applicable to the Checkmark application. Numbers quoted are the percentages of correct detections.

Algorithm (bits) (images, attacks)	Non-geometric (5, 230)	Copyright	Banknote	Logo	Medical	Video
Cox (1)	74%	N/A	N/A	N/A	N/A	N/A
Kim (1)	48%	N/A	N/A	N/A	N/A	N/A
Wang (1)	74%	N/A	N/A	N/A	N/A	N/A
Xia (1)	84%	N/A	N/A	N/A	N/A	N/A
Zhu (1)	93%	N/A	N/A	N/A	N/A	N/A

Table 4. Checkmark 1.2 results for the non-blind algorithms; N/A indicates that data is unavailable, or the algorithm is not applicable to the Checkmark application. Numbers quoted are the percentages of correct detections.

More detailed results for the non-geometric Checkmark application can be found in tables 5 and 6, again, separately for blind and non-blind schemes.

It can be easily seen that the investigated schemes are all relatively robust against image compression; either DCT- or DWT-domain based. However, other attacks pose a serious problem for these technologies, e.g. the maximum likelyhood (ML) estimation attack.

7. CONCLUSION

In this article we have added a number of new attacks to the ones contained in the Checkmark package proposed in prior publications.[19,20] Better understanding of the mechanisms of possible attacks will lead to the development of more efficient and robust watermarking techniques and as such our results present an important step in this direction. We continue to address the issue of application-oriented benchmarking. A myriad of applications have appeared in watermarking and it is now clear that all applications have their own requirements. Work is currently under way to add other applications as well as new attacks. The use of the XML interface greatly facilitates this task and is one of the main advantages of the Checkmark package.

Attack	Fridrich (1)	Dugad (1)	Bruyndonckx (1)	Xie (1)	UniGe2000 (56)
MAP (6)	97%	100%	97%	100%	37%
JPEG (12)	97%	92%	80%	100%	100%
Filtering (3)	93%	80%	100%	100%	100%
Wavelet (10)	86%	98%	54%	100%	62%
ML (7)	43%	91%	60%	100%	26%
Remodulation (4)	15%	100%	0%	100%	10%
Resampling (1)	100%	60%	0%	100%	100%
Copy Attack (1)	0%	100%	20%	100%	60%
Color Reduction (2)	70%	20%	0%	50%	40%

Table 5. Detection performance after non-geometric attack of the blind algorithms; the results are averaged over the five test images

Attack	Wang (1)	Cox (1)	Xia (1)	Kim (1)	Zhu (1)
MAP (6)	80%	100%	100%	30%	100%
JPEG (12)	98%	100%	100%	77%	100%
Filtering (3)	73%	100%	100%	53%	93%
Wavelet (10)	94%	92%	94%	84%	98%
ML (7)	34%	91%	51%	0%	83%
Remodulation (4)	65%	100%	100%	0%	100%
Resampling (1)	40%	100%	70%	0%	80%
Copy Attack (1)	100%	100%	60%	100%	60%
Color Reduction (2)	20%	50%	100%	0%	50%

Table 6. Detection performance after non-geometric attack of the non-blind algorithms; the results were obtained by averaging over the five test images

There is still lot of work left to do. One direction is certainly better integration of perceptual image quality metrics. Another important point is the estimation of the false alarm and false rejection probability. We are looking forward to further improving the Checkmark program.

REFERENCES

1. O. Bruyndonckx, J.-J. Quisquater, and B. M. Macq, "Spatial method for copyright labeling of digital images," in *Proceedings of the IEEE International Workshop on Nonlinear Signal and Image Processing*, pp. 456–459, (Marmaras, Greece), 1995.
2. V. Darmstaedter, J.-F. Delaigle, J.-J. Quisquater, and B. M. Macq, "Low cost spatial watermarking," *Computer & Graphics* **22**(4), pp. 417–424, 1998.
3. I. J. Cox, J. Kilian, T. Leighton, and T. G. Shamoon, "Secure spread spectrum watermarking for multimedia," in *Proceedings of the IEEE International Conference on Image Processing, ICIP '97*, vol. 6, pp. 1673–1687, (Santa Barbara, California, USA), Oct. 1997.
4. J. 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.
5. R. Dugad, K. Ratakonda, and N. Ahuja, "A new wavelet-based scheme for watermarking images," in *Proceedings of the IEEE International Conference on Image Processing, ICIP '98*, (Chicago, IL, USA), Oct. 1998.
6. J. R. Kim and Y. S. Moon, "A robust wavelet-based digital watermark using level-adaptive thresholding," in *Proceedings of the 6th IEEE International Conference on Image Processing, ICIP '99*, p. 202, (Kobe, Japan), Oct. 1999.
7. H.-J. Wang and C.-C. J. Kuo, "Image protection via watermarking on perceptually significant wavelet coefficients," in *Proceedings of the IEEE Workshop on Multimedia Signal Processing*, IEEE, (Los Angeles, CA, USA), Dec. 1998.
8. X.-G. Xia, C. G. Boncelet, and G. R. Arce, "Wavelet transform based watermark for digital images," *Optics Express* **3**, p. 497, Dec. 1998.
9. L. Xie and G. R. Arce, "Joint wavelet compression and authentication watermarking," in *Proceedings of the IEEE International Conference on Image Processing, ICIP '98*, (Chicago, IL, USA), Oct. 1998.
10. W. Zhu, Z. Xiong, and Y.-Q. Zhang, "Multiresolution watermarking for images and video: a unified approach," in *Proceedings of the IEEE International Conference on Image Processing, ICIP '98*, pp. 465–468, (Chicago, IL, USA), Oct. 1998.
11. S. Pereira and T. Pun, "An iterative template matching algorithm using the Chirp-Z transform for digital image watermarking," *Pattern Recognition* **33**, pp. 173–175, Jan. 2000.
12. S. Pereira and T. Pun, "Fast robust template matching for affine resistant image watermarking," in *Proceedings of the 3rd Information Hiding Workshop '99*, A. Pfitzmann, ed., vol. 1768, pp. 199–210, Springer, (Dresden, Germany), Sept. 1999.
13. S. Pereira, S. Voloshynovskiy, and T. Pun, "Optimal transform domain watermark embedding via linear programming," *Signal Processing, Special Issue: Information Theoretic Issues in Digital Watermarking* **81**, pp. 1251–1260, June 2001.

14. P. Loo and N. G. Kingsbury, "Motion estimation based registration of geometrically distorted images for watermark recovery," in *Proceedings of SPIE, Security and Watermarking of Multimedia Contents III*, vol. 4314, (San Jose, CA, USA), Jan. 2001.
15. J.-L. Dugelay and F. A. P. Petitcolas, "Possible counter-attacks against random geometric distortions," in *Proceedings of IS&T/SPIE's 12th Annual Symposium, Electronic Imaging 2000: Security and Watermarking of Multimedia Content II*, P. W. Wong and E. J. Delp, eds., vol. 3971, (San Jose, CA, USA), Jan. 2000.
16. C. I. Podilchuk and W. Zeng, "Perceptual watermarking of still images," in *Proc. Electronic Imaging*, vol. 3016, (San Jose, CA, USA), February 1996.
17. F. A. P. Petitcolas and R. J. Anderson, "Attacks on copyright marking systems," in *2nd International Information Hiding Workshop*, pp. 219–239, (Portland, Oregon, USA), April 1998.
18. F. A. P. Petitcolas and R. J. Anderson, "Evaluation of copyright marking systems," in *IEEE Multimedia Systems (ICMCS'99)*, vol. 1, pp. 574–579, (Florence,Italy), June 1999.
19. S. Pereira, S. Voloshynovskiy, M. Madueo, S. Marchand-Maillet, and T. Pun, "Second generation benchmarking and application oriented evaluation," in *International Information Hiding Workshop III*, (Pittsburgh, PA,USA), April 2001.
20. S. Voloshynovskiy, S. Pereira, V. Iquise, and T. Pun, "Attack modelling: Towards a second generation watermarking benchmark," *Signal Processing* , June 2001.
21. A. B. Watson, "DCT quantization matrices visually optimized for individual images," in *Proc. SPIE:Human vision, Visual Processing and Digital Display IV*, vol. 1913, pp. 202–216, SPIE, 1993.
22. S. Pereira, F. Deguillaume, and T. Pun, "Image watermarking using lapped orthogonal transforms," in *Proceedings of the 11th SPIE Annual Symposium, Electronic Imaging '99, Security and Watermarking of Multimedia Contents*, P. W. Wong and E. J. Delp, eds., vol. 3657, (San Jose, CA, USA), Jan. 1999.
23. P. Bas, J.-M. Chassery, and F. Davoine, "Using the fractal code to watermark images," in *IEEE Int. Conference on Image Processing 98 Proceedings*, vol. 1, Focus Interactive Technology Inc., (Chicago, Illinois, USA), October 1998.
24. S. Bhattacharjee and M. Kutter, "Compression tolerant image authentication," in *IEEE Int. Conference on Image Processing 98 Proceedings*, Focus Interactive Technology Inc., (Chicago, Illinois, USA), October 1998.
25. V. Solachidis, A. Tefas, N. Nikolaidis, S. Tsekeridou, A. Nikolaidis, and I. Pitas, "A benchmarking protocol for watermarking methods," in *Proceedings of the IEEE International Conference on Image Processing, ICIP '01*, (Thessaloniki, Greece), Oct. 2001.