EURASIP Journal on Applied Signal Processing 2002:2, 140–142

A Note on Watermark Development from the Commercial Context

David Hilton

Signum Technologies, Witney, UK
Email: daveh@signumtech.com

Received 6 November 2001

The question is raised as to whether or not the currently developed theories of watermarking provide an adequate model for the handling of digital data in the commercial world.

Keywords and phrases: watermarking, commercial, quality metrics data, modelling, screening.

1. INTRODUCTION

The development of watermarking theories seems to have reached something of an impasse in much the same way as many previous attempts to model real situations using mathematical tools. After several reasonable methods were introduced in the early nineties the development followed the usual route of grabbing a few well-accepted mathematical techniques and attempting to apply them rather than investigating the practical contexts and seeking to model them. Thus the process has attempted to follow the route of reductionist theories where the formalism is well established and the problem is to apply the theory in different contexts. A more appropriate paradigm should perhaps have resembled Darwin's development of understanding of evolution where a myriad contexts are explored to whether they confirm or undermine the proposed theory.

The contrast between the methodology of small commercial enterprises and the academic world in dealing with images is to some extent instructive. The commercial developers start from a strong knowledge base in the practical handling of images at all stages and then seize whatever information they can from the various theorists. Development time is at a premium and so decisions about preferred methods must be taken very quickly. The academic world starts with a broad knowledge of signal processing, cryptography, error correction algorithms, and the like and more or less looks for a situation in which to apply them. The image world is more messy than the audio world because the actual medium of paper and ink is less well defined and thus the theories have struggled more to describe images than audio. The need for some sort of dialectic arising from the commercial and the academic would seem to offer positive benefits.

An important practical issue which arises from this is the consideration of whether there can be a single algorithm which covers the watermarking of digital media, images, and audio, of all qualities and for all purposes. The academic approach would work in favor of a single watermarking algorithm without the pressure experienced in commercial terms to be properly adaptive to different contexts. This note raises issues about the development of watermarking and the extent to which it has confronted the practical problems. It then considers how the type of watermark might be influenced by such considerations.

2. PARADIGMS FOR DEVELOPMENT OF WATERMARKING

The development of JPEG provided an illustrative paradigm for the development of watermarking theories both in its strengths and weaknesses. The starting point was the human visual system based on genuine research into human visual characteristics. Observation showed that the eye could resolve certain frequencies at certain intensities and hence the formalism that developed involved the decomposition of pixel data into frequency data, the frequencies being gathered over fixed 8×8 or 16×16 blocks. A weakness of most of the discussion was the metric that was used to assess the quality of the compression. Assessment was generally carried out by summation of squares of differences, or similar quantity, across a whole image. Unfortunately when customers used the compression with high quality images there were two main types of complaint, neither reflected in the metrics. The first was the appearance of occasional artifacts which would scarcely contribute to the degradation index. The second was the fact that some highlights tended to be diminished. Thus in an image of jewellery the diamonds would have lost their sparkle as a result of the averaging which the cosine transform produced but which again did not figure significantly in the degradation index.

The problems arose partly because the JPEG formalism imposed an implicit model of image data that did not correspond to reality. In particular, the 8 × 8 blocks imposed a structure which could be wildly inappropriate to images. Discontinuities that cut across an 8 × 8 block were particularly awkward. A common feature in images is a contour on one side of which is a smooth area and on the other a noisy area and this is not easily described unless quantization of frequencies is very fine. Clearly much work was carried out on image quality but there was always the feeling that researchers had a neat mathematical model in their minds and were reluctant to abandon it to satisfy a few difficult customers.

A similar scenario seems to afflict the application of cryptography where the most elegant formalism do not correspond well to acceptable practice. Schneier [1] describes his shift of emphasis from the highly developed mathematical algorithm approach to security to an approach where the actual context in its entirety is considered.

3. DEVELOPMENT OF WATERMARKING

The development of watermarking, at least as perceived from an admittedly limited view in the commercial sector, seems to have suffered from some of the same problems as JPEG. This is more so in the case of images than audio because audio is a more accessible medium than paper and ink and fits more easily into the signal processing background of most students. It may also reflect the fact that music has had a mathematical formalism for centuries, including perceptual approximations such as the equal temperament scale as approved by Bach. Two particular aspects seem to have affected development.

3.1. *Quality metrics*

Without a reasonable metric much of the discussion of watermarks from the commercial view is meaningless. There is little point in an academic discussion of the problems of the collusion attack or resolution of ambiguities in rightful ownership if the effect of a single watermark has been undesirable, let alone the subsequent tampering that is proposed. Kutter and Petitcolas [2] list a set of "commonly used pixel based visual distortion metrics" which exhibit all the defects mentioned above. He goes on to point out that more recently perceptual quality metrics have been developed to provide a more useful index of success. There is, however, no detail provided of how to address the likelihood of occurrence of unacceptable artifacts, or loss of important features. There is no discussion, for instance, of how the metrics might need to be modified to meet the needs of screened data where the auxiliary signal has been grouped into regular frequency patterns. There is no doubt that it is difficult to define an acceptable level of artifact and that anisotropic effects complicate matters, but in real commercial situations there is the need for software to run through thousands of images and pick out the unacceptable degradations by an automatic procedure. The issue of quality may seem to be overstated but a brief inspection of the proceedings of the prepress industry, for instance, would reveal otherwise. One only has to note the resources that have been attached to producing optimal screening techniques to avoid moire patterning, or the discussions amongst printers as how best to handle colour profiles and the huge number of problems concerned with the properties of inks, to realize that quality assessed at a sophisticated level is all important.

3.2. *Data modelling*

As a plodding applied mathematician the first process in a problem is usually to model the data in some way. Most problems seem to proceed in roughly two stages. In the first the data can be modelled by some convenient mathematical theory. In the second the presence of difficult exceptions asserts itself. The successful problem solver has to affect some compromise that will not be too time consuming.

In the watermarking literature there is little serious attempt to model image data. One expects to hear of white Gaussian noise being the common assumption but one expects more than an independent pixel model to reflect the nature of the underlying data.

Here at Signum, for instance, we discovered that the retrieval of Signum's watermarks varied quite considerably according to the assumptions made about local correlation between pixel values. There was a sharp improvement in performance when a simple, if rough, method of computing relevant local variance was introduced.

In the same vein Herrigel et al. [3] made important observations about the need for localized noise analysis where they described a noise visibility function (NVF) which was important in the robustness of the watermark. In the formula they derived the NVF which turns out to be inversely proportional to the local energy defined by the local variance. One would imagine that this could be taken further by considering adaptations needed for a variety of possible image features and where data is sharply discontinuous. A model of the behaviour of screened data where the energy is grouped periodically might provide further refinement.

4. WATERMARKING STYLES

The above suggests that there may be some benefit in reviewing types of watermark in the light of better metrics and better data modelling and hence coming to a clearer decision about the possibility of a universal algorithm. There is an intuitive feeling that certain issues, exemplified below, might need more consideration.

As the foregoing discussion implies an essential requirement is that a watermark intensity should be adaptable to the local data. It would seem that this would give a low rating to schemes that have a predetermined form, as for instance, schemes that modify frequencies over large areas of an image, knowing that in some cases the texture of the data will vary violently over small distances. The correlation required for detection implies having access to a coherent subset of the image and this access may be seriously diminished by the type of transformation brought about by cropping and printing.

In contrast, methods such as the "Patchwork" method of Bender et al. [4] which accumulate data from chosen points and do not require any geometrical extent, might be more robust. The Patchwork method makes it easy to vary the intensity of a watermark across an image and yet still ensure that every pixel makes a contribution to the overall detection. This local responsiveness must be a positive for a watermark from the perceptibility viewpoint and was, in fact, part of the reason for Signum to select that type of signal. This together with a permutation scheme that distributes the signal in a more secure manner and leads to a more rapid searching algorithm produces a method which has many practical advantages. Schemes based on wavelet analysis would appear to offer the prospect of more adaptable watermarks providing as they can orthogonal functions that model well the significant parts of the data.

The modification of frequencies within images must be regarded with some degree of scepticism by anyone concerned with quality. Everyone, I imagine, accepts Cox's [5] assertion of the need for the watermark to be substantially present in the perceptible parts of the data on the grounds that any compression scheme or tolerable image transformation may well attenuate other parts. One can take this further and assert that a watermark should be present up to a level where it is just below perceptibility so that any attempt to add further watermarks will result in a degraded image. However, images may be subject to spurious frequency addition in the course of screening, scanning and JPEG compression and the interaction of the processes with DCT modifications, for instance, may well be damaging.

These are perhaps hand waving arguments about information embedding but in the absence of detailed data models it is difficult to provide more. I look to the academic community to push their research a little nearer to genuine mathematical modelling.

REFERENCES

[1] B. Schneier, *Applied Cryptography*, John Wiley and Sons, New York, NY, USA, 2nd edition, 1996.

[2] M. Kutter and F. A. P. Petitcolas, "A fair benchmark for image watermarking systems," in *Proc. SPIE Security and Watermarking of Multimedia Contents*, P. W. Wong and E. J. Delp, Eds., vol. 3657, pp. 25–27, San Jose, Calif, USA, January 1999.

[3] S. Voloshynovskiy, A. Herrigel, F. Jordan, N. Baumgärtner, and T. Pun, "A noise removal attack for watermarked images," in *Multimedia and Security Workshop*, J. Dittmann, K. Nahrstedt, and P. Wohlmacher, Eds., Orlando, Fla, USA, October 1999, (at the 7th ACM Multimedia Conference (Multimedia 99)).

[4] W. Bender, N. Morimoto, and D. Gruhl, "Method and apparatus for data hiding in images," United States Patent 5,870,499, February 1999.

[5] I. J. Cox, J. Kilian, T. Leighton, and T. Shamoon, "Secure spread spectrum watermarking for multimedia," *IEEE Transactions on Image Processing*, vol. 6, no. 12, pp. 1673–1687, 1997.

David Hilton graduated in Mathematics at Oxford University and gained his Ph.D. in Theoretical Physics. He spent the first part of his career in Education in the course of which he carried out a wide range of mathematical consultancy including practical security models. In 1990, he transferred full time to industry carrying out work on image compression, colour manipulation and quality issues. In 1995, he became involved in watermarking for both images and audio with Signum Technologies.