Digital watermarking using improved human visual system model

Brett T. Hannigan, Alastair Reed, Brett Bradley

Digimarc Corporation, 19801 SW 72nd Ave., Suite 250, Tualatin, OR 97062

ABSTRACT

In digital watermarking, one aim is to insert the maximum possible watermark signal without significantly affecting image quality. Advantage can be taken of the masking effect of the eye to increase the signal strength in busy or high contrast image areas. The application of such a human visual system model to watermarking has been proposed by several authors. However if a simple contrast measurement is used, an objectionable ringing effect may become visible on connected directional edges.

In this paper we describe a method which distinguishes between connected directional edges and high frequency textured areas, which have no preferred edge direction. The watermark gain on connected directional edges is suppressed, while the gain in high contrast textures is increased. Overall, such a procedure accommodates a more robust watermark for the same level of visual degradation because the watermark is attenuated where it is truly objectionable, and enhanced where it is not.

Furthermore, some authors propose that the magnitude of a signal which can be imperceptibly placed in the presence of a reference signal can be described by a non-linear mapping of magnitude to local contrast. In this paper we derive a mapping function experimentally by determining the point of just noticeable difference between a reference image and a reference image with watermark.

Keywords: Digital watermarking, Edge Detection, Connected Edges, Human Visual System, Digimarc

1. INTRODUCTION

The process of digital watermarking involves tweaking pixel values at various regions within an image to encode some piece of information. This encoded information can be used to provide copyright information, to prevent illegal duplication, or even as a dynamic link between the image and online digital data. For most applications, the image owner would like to make the encoded data robust enough to ensure its detection while maintaining the high quality of the original image. Often times these two desires are in direct conflict with each other – any increase in robustness comes at the expense of increased signal visibility. However, the visual impact a signal has on a particular image is not solely related to the magnitude of the signal, nor is the detectability of a signal solely related to its magnitude. Local image characteristics also help determine the visibility of a signal and its detectability. For instance, a signal embedded into an image area which is primarily flat will be more visible than a signal of equal magnitude embedded into an area of high contrast. This latter signal may be more difficult to detect however, as areas of higher contrast can interfere or jam a signal more easily than a flat area.

Many watermarking schemes attempt to analyze images in order to identify how much watermark signal various image areas can hold while maintaining equal visibility. A related methodology attempts to identify how much watermark signal various image areas require to maintain equal detectability across the entire image. These techniques often accomplish this by calculating a value of local contrast, and mapping increasing contrast values to increasing signal magnitudes[1],[2]. This can present a problem though, as directional edges separating two distinct objects in an image may be identified as high contrast areas. This results in the application of a higher strength watermark signal around the connected edge, which causes an objectionable watermark ringing or halo.

In this paper, we introduce a method which better identifies areas of true high contrast texture while protecting connected directional edges. Additionally, we describe a more accurate mapping to be used when mapping local contrast to watermark gain. We then discuss results of utilizing the new algorithm, and compare those results to a previous method. We conclude by discussing possible extensions to this method for the future.

2. ALGORITHM DESCRIPTION

Our algorithm for determining the magnitude of watermark signal to be applied based on local image characteristics can be described in four conceptual stages, as shown in **figure 1**.

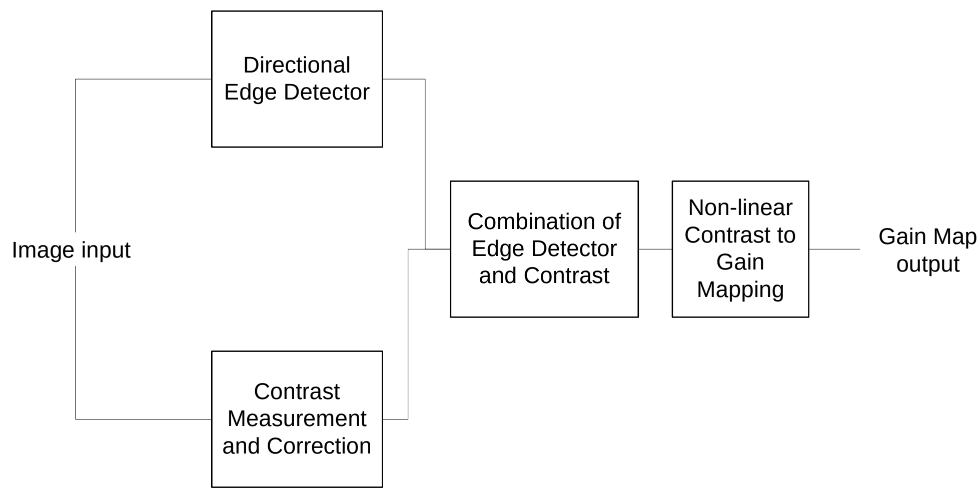

Figure 1. Watermark Gain Algorithm Block Diagram

The image data is first fed into two separate stages – the contrast measurement stage and the directional edge detector stage. The results from these two stages are then combined to form a "corrected contrast measurement map" which reports the local contrast values for the image while protecting directional edges. The results from this stage are then passed into a non-linear contrast to gain mapping stage which calculates the gain values, or magnitude of watermark signal varied image regions should receive.

2.1 Contrast Measurement

If a watermark is applied at equal strength throughout an image it will tend to be more visible in texturally flat regions, and less visible in busier areas. Conversely, it is usually more difficult to extract the watermark from pixels in busy regions than in those that are texturally flat. For these two reasons it is in general desirable to measure the textural contrast of the image to be watermarked on a local basis. The obtained measurement is then used to control the strength of the watermark applied to the measured region. The process would be repeated for all regions of the image that require a watermark.

The measurement of contrast typically involves one or more filtering operations, possibly non-linear. To make an initial measurement of contrast, our implementation uses a band pass filtering operation. Although straightforward filtering produces a good initial result, refinement is required before the textural contrast measurement can be mapped to watermark strength without undue ill effects.

Using filtering to determine textural contrast, and hence watermark strength, without taking into account certain natural image characteristics leads to apocryphal indicators of how heavily a watermark should be applied to a given region. A region that contains a sudden transition in luminance may be labeled as a prime candidate for high watermark strength after initial contrast filtering. Regions that contain borders, text, or fine lines are some examples. If the region is heavy-handedly doused with a watermark it may appear objectionable depending upon the characteristics of the region and others that surround it. We classify as false contrast regions those that cannot truly support high watermark strength when the filter-based contrast measurement would indicate otherwise.

One method of dealing with potentially false contrast regions is to de-emphasize, or even penalize, such regions if they have an uncharacteristically high contrast. For example, we have been able to characterize our contrast filter by applying it to images we would regard as busy; the image does not degrade noticeably under high watermark strength. We found that on average the contrast measurement is relatively low compared with many of the false contrast regions. By characterizing our filter we set a peak expected contrast. If a region's contrast supersedes the expected peak contrast, its final assigned contrast is clipped at the expected peak value, or in some case reduced below the peak value. Although beneficial, the described contrast adjustment procedure works only on a local image region basis. By taking into account groups of regions, more intelligent decisions can be made regarding the application of watermark strength. Our so-called directional edge finding method serves to do just that.

2.2 Connected Edge Detection

Edge detection algorithms have been evaluated in image processing literature [3],[4]. Typical edge detection processes define an edge as a "step discontinuity in the image signal" [5], and attempt to locate edges by convolving the image with a kernel that approximates a first or second derivative. Using a first derivative kernel, an edge occurs at local maxima, while for second derivatives, edges occur at the zero crossings. John F. Canny developed the current standard in edge detection [6]. His method starts by convolving an image with a 2-D Gaussian filter, and then differentiating this smoothed image in two orthogonal directions. By calculating the derivative in two orthogonal directions, one can determine the overall gradient direction and amplitude. Using this knowledge, the algorithm then suppresses points which were non-maxima and values that were not local peaks. The final step involves thresholding the edges. Two threshold values are used . The first threshold, which is larger than the second, identifies most of the true edges. However, some discontinuities in edges may occur using this higher value, so pixels which are connected to these high threshold edges are also considered edges if they are above the second, smaller threshold. The resulting map typically provides a very good representation of the edges in the image.

The edge map provided by the Canny edge detection algorithm does not attempt to differentiate between connected edges and areas with a high concentration of random edge like textures. For example, for an image like "Mandrill", a Canny edge algorithm might construct the edge map in **figure 2a**, in which the fur around the mandrill's nose is said to contain a high concentration of edges. However, the purpose of the edge map in our algorithm is to highlight connected edges we should avoid when increasing watermark signal gain. The fur area should be able to hold a good deal of watermark signal, since the so-called edges are really a somewhat random texture, and placing a noise like signal in this texture is unlikely to be noticed. Thus, we modify Canny's edge detection algorithm to ignore random, closely packed edges in the following fashion. First, we take the edge map provided by Canny's algorithm and smear it with a 5x5 or 7x7 low pass filter kernel. This causes closely packed edges to bleed into one another. Next, we thin this smeared edge map using a min filter of a slightly higher order than the smear. This causes edges which were stretched to contract back in tighter than the original Canny edge mask. A composite, binary edge mask is then constructed by saying a pixel is on an edge if and only if the original Canny edge mask says it is an edge, but the min-filtered, smeared Canny Edge mask says it is not. This operation essentially allows edges that are boundary edges to remain, while closely packed edges disappear. The final step grows the edge map by a suitable radius to protect all pixels which may inadvertently be called "high contrast" areas due to these edges. The effect of this process on "Mandrill", as can be seen in **figure 2b,** is to keep the edges around the mandrill's eyes, nose, and outer face, while disregarding the texture edges on his fur.

Figure 2a. Canny Edge Mask of "Mandrill"

Figure 2b. Our modified Canny Edge Mask of "Mandrill"

2.3 Combination of Edge Detection and Contrast Measurement

The third stage of our algorithm combines the results of the edge detection and contrast measurement stages. As stated above, the general method in determining the amount of watermark gain a certain image area should receive is mapping local contrast to some gain value. The purpose of the edge detection section is to identify areas which may be perceived as high contrast areas by the contrast algorithm, but in reality can not hold the watermark gain generally associated with such a high contrast. Examples of these regions are text edges, object boundaries, and other directional edges.

This combination stage therefore takes as input the contrast map of the image as well as the binary edge map. For any contrast value which is not said to be on an edge, the contrast value is untouched. For those contrast values which are found to be on a directional edge however, the reported contrast value is calculated as a percentage of the original. We have found experimentally that a contrast reduction of 50% to 80% on edges provides a clean gain map with edges free of objectionable watermark ringing.

2.4 Non-Linear Contrast to Gain Mapping

The combined detail gain contrast calculation is then passed through a one-dimensional mapping function shown below in **figure 3** to obtain a final detail gain.

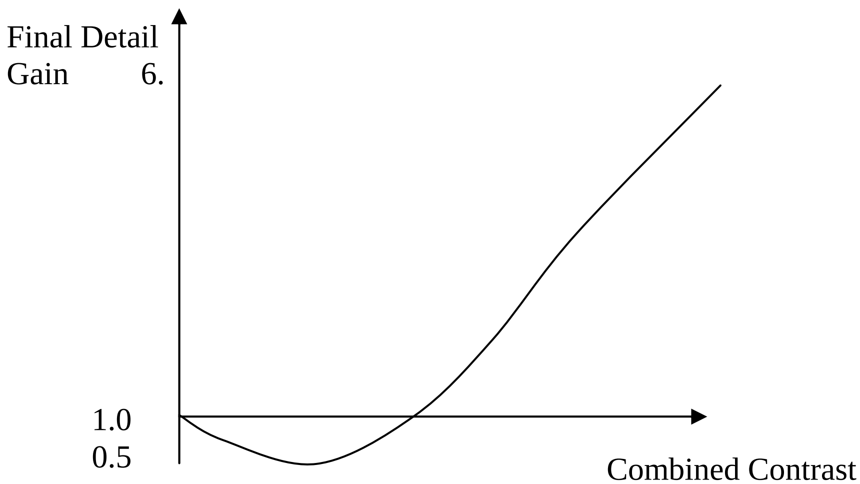

Figure 3. Combined Contrast to Watermark Gain Mapping Function

The detail gain mapping function has a dip at low contrast, and is then approximately linear on a log scale. The function was calibrated for our application by using a mid-gray patch with white noise of various contrast levels, and embedding at different strengths until the watermark is just perceptually visible in order to build up the curve shape. Further tests were then performed on a standard image set. The shape obtained was very similar to the generalized contrast discrimination model reported in [7]. The model shows that the peak sensitivity of the human visual system to a signal, in the presence of a reference signal, is at a low contrast of about 0.9 % modulation of the reference signal.

3. RESULTS

3.1 Comparison of Two Watermark Gain Methods

In **figure 4a**, watermark gain was simply varied depending on local image contrast For some connected edges, as in the road sign, this resulted in a visually objectionable watermark that appeared as a ringing on connected edges. The proposed watermark gain with edge suppression method, identifies these areas as connected edges, and suppresses the watermark signal in these regions. The result shown in **figure 4b** is an image with much cleaner text edges, since the objectionable ringing on connected edges has been suppressed.

Figure 4a. Watermark Gain Based on Simple Contrast (watermark accentuated for effect.)

Figure 4b. Watermark Gain Based on Contrast with Edge Suppression (watermark accentuated for effect.)

The non-linear curve that is used to apply the watermark gain based on contrast takes advantage of the masking effect of the human eye, allowing a higher strength signal to be put into high contrast areas. This could not be done in the past due to the objectionable ringing of the watermark on connected edges. Since the watermark gain is suppressed on connected edges where the visibility is most objectionable, the gain can be increased to a higher level in other busy areas that are able to support a high watermark signal without visibility problems. Thus a reduction in overall watermark visibility is obtained at about the same detection rate.

The detection rate is not greatly affected by the watermark suppression on connected edges, since a very small fraction of the overall number of pixels in the image are reduced in strength. Also, in many watermarking schemes edge regions are much more difficult to read in general. Thus the overall detection rate is not greatly affected by the watermark suppression on connected edges.

Figure 5a shows the standard mandrill image. **Figure 5b** shows the image after a watermark has been applied using the new gain algorithm. A strong watermark signal has been applied without any significant visual problems being apparent.

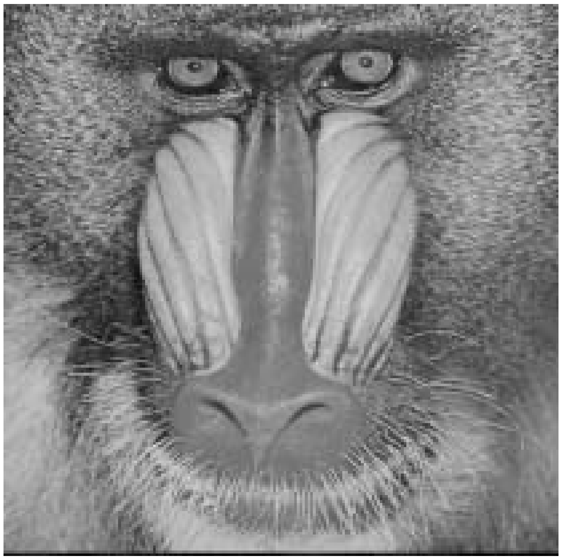

Figure 5a. Mandrill without Watermark

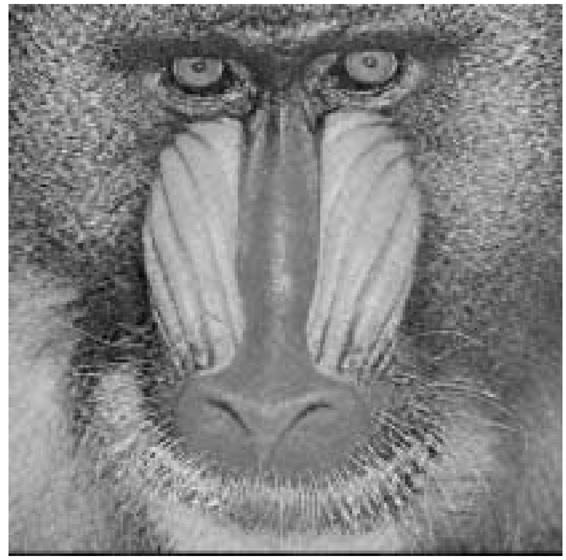

Figure 5b. Mandrill with Watermark using New Gain Algorithm

3.2 Algorithm Limitations

While the algorithm developed in this paper succeeds in protecting most connected edges and false-high-contrast areas from receiving too much watermark gain, it does have some limitations. Due to the nature of the modified Canny edge detection algorithm, areas with narrow, parallel lines will still be considered true-high-contrast areas. This will cause image areas such as the girl's hair in the image labeled **figure 6b** to receive a large watermark gain, which appears visually objectionable. Additionally, small, closely spaced text may still have areas which receive too much watermark gain, although not nearly as much as in a simple contrast method. This is due to the fact that in the modified Canny edge detection algorithm, the smeared Canny edge map may smear text together, and the min filter will not shrink this area back inside the original text boundaries. This would leave some edges of the text with uncorrected high-contrast values.

Figure 6a. Image of Child without Watermark (watermark accentuated for effect.)

Figure 6b. Image of Child with Watermark using New Gain Algorithm (watermark accentuated for effect.)

4. EXTENTIONS AND CONCLUSION

The algorithm developed in this paper successfully uses a local contrast value to determine the watermark gain, while additionally suppressing gain at most false-high-contrast areas. In certain instances, the algorithm mistakenly allows a large watermark gain in an area that is unable to imperceptibly hold it. Future versions of the algorithm should therefore work to better differentiate between random, high contrast texture, and more uniform, parallel line patterns, as well as closely spaced text.

This procedure has been included in Digimarc's MediaMarc Embedder v1.5 product, which can be used to embed digital watermarks into professional print products, such as magazines, newspaper, catalogs, and even CD liner notes. For these applications, visual quality is of the utmost concern, so the enhancements in image quality afforded by a more intelligent watermarking algorithm are greatly appreciated. Generally speaking, as watermarking becomes less visually perceptible, watermarking as a copyright information system, copy-prevention system, or content enabling system will be more widely accepted.

ACKNOWLEDGEMENTS

The authors would like to thank Tony Rodriguez, Don Haaga, Lou Berkley, Eric Hudson, Steve Whaley, Matt Riley, Peter Harmon, Ben Bounketh, David Cushman, Brian MacIntosh, Steve Decker, Burt Perry, Geoff Rhoads and Joel Meyer of Digimarc for their help in contributions to this paper.

REFERENCES

1. C.I. Podilchuk, W. Zeng, "Image-adaptive watermarking using visual models," 1998.
2. J.F. Delaigle, C. De Vleeschowuwer, B. Macq, "Watermarking algorithm based on a human visual model," *Signal Processing*. **66**, pp. 319-335, 1998.
3. J. S. Lim, *Two-Dimensional Signal and Image Processing*, pp. 476-495, PTR Prentice Hall, New Jersey, 1990.
4. W. K. Pratt, *Digital Image Processing*, pp. 478-492, John Wiley & Sons, New York, 1978.
5. P. Kovesi, "Lecture 6, Classical Feature Detection," http://www.cs.uwa.edu.au/undergraduate/courses/233.412/Lectures/Lecture6/lecture6.html, The University of Western Australia, 2000.
6. J.F. Canny, "A computational approach to edge detection," *IEEE Trans. Pattern Analysis and Machine Intelligence*, **8, pp.** 679-698, 1986.
7. Peter G.J. Bartens, ‘Contrast Sensitivity of the Human Eye and Its Effects on Image Quality’, p.147-151, SPIE Press, 1999