

A Vision-Based Masking Model for Spread-Spectrum Image Watermarking

Martin Kutter[†] and Stefan Winkler[‡]

Signal Processing Laboratory
Swiss Federal Institute of Technology
1015 Lausanne, Switzerland
<http://ltswww.epfl.ch/>

Abstract— We present a perceptual model for hiding a spread-spectrum watermark of variable amplitude and density in an image. The model takes into account the sensitivity and masking behavior of the human visual system by means of a local isotropic contrast measure and a masking model. We compare the insertion of this watermark in luminance images and in the blue channel of color images. We also evaluate the robustness of such a watermark with respect to its embedding density. Our results show that this approach facilitates the insertion of a more robust watermark while preserving the visual quality of the original. Furthermore, we demonstrate that the maximum watermark density generally does not provide the best detection performance.

Keywords— Isotropic contrast, masking model, spread-spectrum watermark, robustness, density.

I. INTRODUCTION

The ease with which images in digital form can be distributed and reproduced has become a growing concern for their creators, and protecting authors' rights has become increasingly important. One solution to these problems is digital watermarking, i.e. the insertion of information into the image data in such a way that the added information is not visible and yet resistant to image alterations. A variety of techniques has already been proposed; an overview of the subject can be found in [8].

That digital watermarks should be invisible prompts the exploitation of the characteristics of the human visual system (HVS). In general, this implies a spatial or spectral shaping of the watermark according to certain HVS (in)sensitivities. However, this idea has not been pursued vigorously until recently. Kutter et al. [11] proposed modifying the blue component of the image proportional to the local luminance. More sophisticated approaches were presented by Podilchuk and Zeng [17], who insert a watermark in the DCT domain, or by Delaigle et al. [4], who consider the masking of band-limited noise in texture regions and around contours.

In this paper, we propose a new HVS-optimized weighting function for hiding a spread-spectrum watermark in an image. Both luminance and blue-channel watermarks are investigated. We use a weighting function derived from the masking behavior of the human visual system as well as from the watermark itself. Our results show that this approach facilitates the insertion of a more robust watermark in either the luminance image or in the blue channel while preserving the visual quality of the original.

The paper is organized as follows: Spread-spectrum watermarking is introduced in Section II. The characteristics of the human visual system used in our approach are presented in Section III. Application-specific design choices for our system are described in Section IV, and the experiments carried

out to determine the pertinent model parameters are reported in Section V. Finally, we evaluate the proposed watermarking schemes with respect to their visibility and robustness and discuss our results in Section VI.

II. SPREAD-SPECTRUM WATERMARKING

The approach chosen for the digital watermark is based on spread-spectrum modulation [16, 19]. In this section we present the watermark insertion and detection process [9, 10].

A. Watermark Insertion

A watermarked image $\hat{I}(x, y)$ is created by adding a watermark $w(x, y)$ to the original image $I(x, y)$:

$$\hat{I}(x, y) = I(x, y) + w(x, y). \quad (1)$$

The watermark $w(x, y)$ contains an N -bit binary signature $B = \{b_0, \dots, b_{N-1}\}$. For each bit i , a two-dimensional modulation function $s_i(x, y)$ is generated as follows: we start with a set of pixel positions S_i spanning the image. To obtain orthogonality between the modulation functions, these sets are defined such that their intersection is empty, i.e. $S_i \cap S_j = \emptyset \quad \forall i \neq j$. This condition ensures that each pixel of the image is modified by only one function. The modulation functions are then defined as follows:

$$s_i(x, y) = \begin{cases} p_i(x, y) & \text{if } (x, y) \in S_i, \\ 0 & \text{otherwise.} \end{cases} \quad (2)$$

The functions $p_i(x, y)$ are pseudo-random functions. In our case, they have a bimodal distribution of $\{-1, 1\}$, but other distributions are viable.

The watermark itself is defined as the linear superposition of N modulated and weighted functions $s_i(x, y)$:

$$w(x, y) = \sum_{i=0}^{N-1} (-1)^{b_i} s_i(x, y) \alpha(x, y), \quad (3)$$

where $\alpha(x, y)$ represents a weighting function, which will be defined in Section IV.

To gain greater control over the artifacts, we introduce a watermark density D , which defines the fraction of pixels in the image to be watermarked:

$$D = \frac{|\{\bigcup_{i=0}^{N-1} S_i\}|}{|\{S\}|}, \quad (4)$$

where $|\{\cdot\}|$ is the cardinal of a set, and $\{S\}$ is the set of all positions in the image. The spatial probability distribution of the positions in a set $s_i(x, y)$ is uniform, meaning that the probability for a pixel to be part of set S_i is D/N .

[†] MK is now with AlpVision, Switzerland [Martin.Kutter@alpvision.com]

[‡] SW is now with Genimedia, Switzerland [swinkler@genimedia.com]

B. Watermark Detection

In order to demodulate the information inserted in the image, the correlation between the watermarked image and the modulation functions is computed by means of a linear correlator. Since the statistical properties of an image are not stationary and the expectation is non-zero, the watermarked image is pre-processed in order to reduce its variance. This procedure greatly improves the performance of the watermarking system. The detector statistics are given by:

$$r_i = \langle f(\hat{I}), s_i \rangle, \quad (5)$$

where f is the pre-processing function, \hat{I} is the watermarked image, and s_i are the modulation functions.

We employ adaptive watermark prediction based on a Wiener filter. Wiener filtering is commonly used in image restoration and denoising, because it exhibits optimal performance when both image and noise statistics are Gaussian. The coefficients $h_0(m, n)$ of an adaptive Wiener filter of size $W \times W$ are defined as follows [12]:

$$h_0(m, n) = \begin{cases} \frac{1}{\sigma_I^2 + \sigma_w^2} \left(\sigma_I^2 + \frac{\sigma_w^2}{W^2} \right) & \text{if } m = n = 0, \\ \frac{1}{\sigma_I^2 + \sigma_w^2} \frac{\sigma_w^2}{W^2} & \text{if } 0 < |m|, |n| < \frac{W}{2}, \\ 0 & \text{otherwise,} \end{cases} \quad (6)$$

where σ_I^2 is the local variance of the original image and σ_w^2 is the local variance of the watermark.

The watermark variance is computed using an estimate $\tilde{\alpha}$ of the local watermark weighting function $\alpha(x, y)$. Taking into account that the watermark is not present at every point due to the variable embedding density D introduced in Section II-A, the local estimate of the embedded watermark is given by:

$$\tilde{\sigma}_w^2(m, n) = \frac{D}{W^2} \sum_{k=m-W_1}^{m+W_1} \sum_{l=n-W_1}^{n+W_1} \tilde{\alpha}^2(k, l), \quad (7)$$

where $W_1 = \frac{W-1}{2}$.

Under the assumption that the variance of the image is independent of the watermark noise, its estimate is given by:

$$\tilde{\sigma}_I^2(m, n) = \max \{ \tilde{\sigma}_I^2(m, n) - \tilde{\sigma}_w^2(m, n), 0 \}, \quad (8)$$

where $\tilde{\sigma}_I^2(m, n)$ is an estimate of the local variance of the watermarked image:

$$\tilde{\sigma}_I^2(m, n) = \frac{1}{W^2} \sum_{k=m-W_1}^{m+W_1} \sum_{l=n-W_1}^{n+W_1} (\hat{I}(k, l) - \tilde{\mu}_I(m, n))^2, \quad (9)$$

and $\tilde{\mu}_I(m, n)$ is an estimate of the local mean:

$$\tilde{\mu}_I(m, n) = \frac{1}{W^2} \sum_{k=m-W_1}^{m+W_1} \sum_{l=n-W_1}^{n+W_1} \hat{I}(k, l). \quad (10)$$

The Wiener filter h_0 computes an estimate of the original image. In the watermark detection process, the difference between the watermarked image and this estimated original is used. This difference can be included directly in the adaptive Wiener filter as follows:

$$h(m, n) = \begin{cases} \frac{W^2-1}{\sigma_I^2 + \sigma_w^2} \frac{\sigma_w^2}{W^2} & \text{if } m = n = 0, \\ -\frac{1}{\sigma_I^2 + \sigma_w^2} \frac{\sigma_w^2}{W^2} & \text{if } 0 < |m|, |n| < \frac{W}{2}, \\ 0 & \text{otherwise.} \end{cases} \quad (11)$$

Thus pre-processing consists of a convolution with this filter, and the detector statistics from Eq. (5) simply become

$$r_i = \langle h * \hat{I}, s_i \rangle. \quad (12)$$

The signature bits of the watermark can now be determined from the signs of r_i :

$$\tilde{b}_i = \begin{cases} 1 & \text{if } r_i \geq 0, \\ 0 & \text{if } r_i < 0. \end{cases} \quad (13)$$

III. MODEL OF THE HUMAN VISUAL SYSTEM

To hide the watermark in an image effectively, it is useful to exploit certain characteristics of the human visual system. We focus on the masking phenomenon in this paper, because it describes interactions between stimuli such as the watermark and the image. Masking occurs when a stimulus that is visible by itself cannot be detected due to the presence of another. The amount of masking depends on the contrast of the masker. The following two sections discuss a suitable measure of contrast and a masking model for spread-spectrum watermarking.

A. Contrast

The response of the human visual system depends much less on the absolute luminance than on the relation of its local variations to the surrounding luminance. This property is known as *Weber's law*. Contrast is a measure of this relative variation of luminance. Unfortunately, a common definition of contrast suitable for all situations does not exist. Mathematically, Weber contrast can be expressed as:

$$C^W = \frac{\Delta L}{L}. \quad (14)$$

It is often used for small patches with a luminance offset ΔL on a uniform background of luminance L . In the case of sinusoids or other periodic patterns with symmetrical deviations ranging from L_{\min} to L_{\max} , which are also very popular in vision experiments, Michelson contrast is generally used:

$$C^M = \frac{L_{\max} - L_{\min}}{L_{\max} + L_{\min}}. \quad (15)$$

While these definitions are good predictors of perceived contrast for simple stimuli, they fail when stimuli become more complex [15]. It is evident that they are not appropriate for measuring contrast in natural images, because a few very bright or very dark points would determine the contrast of the whole image. Actually, human contrast sensitivity varies with the *local* average luminance. In order to address these issues, Peli [14] proposed a local band-limited contrast measure:

$$C_j^P(x, y) = \frac{\psi_j * I(x, y)}{\phi_j * I(x, y)}, \quad (16)$$

where ψ_j is a band-pass filter at level j of a filter bank, and ϕ_j is the corresponding low-pass filter. Modifications of this local band-limited contrast definition have been used in a number of vision models [2, 13] and are in good agreement with psychophysical experiments on Gabor patches [15].

However, real (symmetric) filters alone are not sufficient for describing the contrast of complex stimuli such as natural images. This becomes obvious when C^P is computed for sinusoids with a constant C^M : The contrast measured only with real (symmetric) filters actually oscillates with the same frequency as the underlying sinusoid [22]. This complicates the establishment of a correspondence between such a local contrast measure and data from psychophysical experiments.

To relate a local definition of contrast to the Michelson contrast of a sinusoidal grating (i.e. a constant response across the whole image), the contrast definition must take into account both the in-phase and the quadrature component [3]. Analytic filters represent an elegant way to achieve this. However, general two-dimensional analytic filters are not well-defined because of the lack of a Hilbert transform in two dimensions [20]. *Directional* analytic filters pose no difficulties, as long as their angular support is smaller than π , allowing meaningful oriented contrast measures to be computed easily [13]. For the application at hand, an isotropic contrast measure is required, since a spread-spectrum watermark exhibits no prominent orientation itself; however, isotropic analytic filters as such do not exist.

We therefore rely on a method proposed recently by Winkler and Vandergheynst [22]. They show how to construct an *isotropic* local contrast measure C^I from the combined responses of analytic directional band-pass filters ψ_{jk} (level j , orientation k). This isotropic local contrast is defined as follows:

$$C_j^I(x, y) = \frac{\sqrt{2 \sum_k |\psi_{jk} * I(x, y)|^2}}{\phi_j * I(x, y)}. \quad (17)$$

The resulting isotropic contrast for level j corresponds to the square root of the energy sum over all orientations of accordingly band-pass filtered images, normalized by the low-pass (ϕ_j) filtered image. The filters ψ_{jk} must be designed such that their angular support is smaller than π and their energy sum over all filter orientations is isotropic in the frequency domain. Then the L^2 -norm of the numerator is equivalent to what would have been obtained using an isotropic filter [22]. C^I thus behaves as prescribed with respect to sinusoidal gratings (i.e. $C^I(x, y) \equiv C^M$ in this case).

Examples of this isotropic contrast applied to the *lena* image at three different levels are shown in Figure 1. The figures clearly illustrate that C^I exhibits the desired omnidirectional and phase-independent properties. Its localization depends on the chosen level. The specific filters used for the contrast computation are discussed in Section IV.

B. Masking

In the context of watermarking it is helpful to think of the watermark being masked by the original image. Masking explains why the watermark signal is disturbing in certain regions of an image while it is hardly noticeable elsewhere.

Masking effects are usually quantified by means of detection experiments, where the contrast threshold for the detection of a target stimulus is measured for a range of masker contrasts. Figure 2(a) shows an example of curves approximating the data typically resulting from such experiments. The horizontal axis shows the log of the masker contrast C_M , and the vertical axis shows the log of the target contrast C_T at detection threshold. For contrast values of the masker larger than C_{M_0} , the detection threshold grows as a power of the masker contrast. The detection threshold for the target stimulus without any masker is indicated by C_{T_0} . C_{M_0} depends on the nature of the stimuli, as we will see in Section V.

Two cases can be distinguished in Figure 2(a). In case A, there is a gradual transition from the threshold range to the masking range. Typically this occurs when masker and target have different orientations or otherwise different characteristics. For case B, the detection threshold for the target *decreases* when the masker contrast is close to C_{M_0} , which implies that the target is easier to perceive due to the presence of the masker in this contrast range. This effect is known as *facilitation* and occurs mainly when target and masker have very similar properties.




Fig. 1. Three levels of isotropic contrast $C_j^I(x, y)$ for the *lena* image.

Masking is generally strongest when the interacting stimuli have similar characteristics, i.e. similar frequencies, orientations, and colors. Masking between stimuli from different channels is weaker [5, 6]. Color masking and interactions between color and luminance stimuli were investigated by Switkes et al. [21]. Their experiments showed that the overall behavior is similar to luminance masking.

Because the facilitation effect is usually limited to a rather small range of masker contrasts, and because in our application the masker (the image) and the target (the watermark) have different characteristics, facilitation is neglected here. The most important effect is masking, i.e. the significant increase of the target's visibility threshold with increasing masker contrast C_M . Hence a simplified masking model can be formulated as follows:

$$C_T(C_M) = \begin{cases} C_{T_0} & \text{if } C_M < C_{M_0}, \\ C_{T_0} (C_M/C_{M_0})^\epsilon & \text{otherwise.} \end{cases} \quad (18)$$

This model is illustrated in Figure 2(b). The masking curve is divided into a threshold range, where the target detection threshold is independent of the masker contrast, and a masking range, where it grows with the power of the masker contrast. The model comprises three parameters, namely ϵ , C_{T_0} and C_{M_0} , which will be determined in Section V by means of subjective experiments.

IV. APPLICATION TO WATERMARKING

A. Lightness Function

Weber's law generally overestimates the sensitivity to high luminance values, leading to a higher visibility of the watermark in bright areas of the image. To compensate for this effect and to have constant watermark visibility over all image intensities,

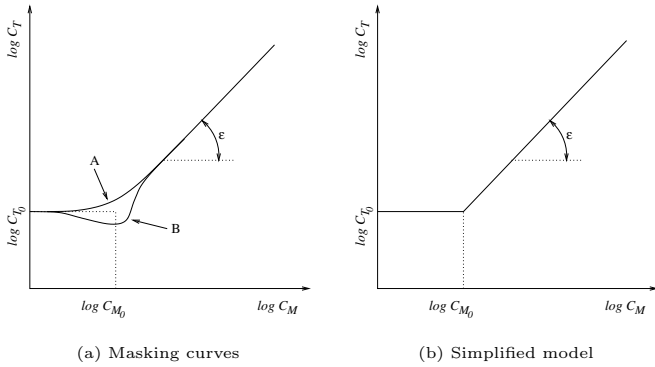


Fig. 2. Illustration of typical masking curves (a) and a commonly used model (b) as defined by Eq. (18). For stimuli with different characteristics, masking is the dominant effect (case A). Facilitation occurs for stimuli with similar characteristics (case B).

a lightness function is applied to the image before contrast computation. We use the modified log characteristic proposed by Schreiber [18]:

$$B(L) = 1 + 99 \frac{\log(1 + L \cdot a) - \log(1 + a)}{\log(1 + 100a) - \log(1 + a)}, \quad (19)$$

where L is the luminance. This characteristic is very flexible. For small a the relationship is approximately logarithmic, while for large a it is almost linear. Heuristic tests by Schreiber indicate that $a = 0.05$ provides a well-adapted luminance scale. This value has been verified in our watermarking application for both embedding channels.

B. Filter Design

For the computation of the local contrast according to Eq. (17) we use directional wavelet frames as described in [1] based on the *PLog* wavelet:

$$\psi_\tau(\vec{x}) = \frac{1}{\tau} \tilde{\psi}_\tau \left(\frac{\vec{x}}{\sqrt{\tau}} \right), \quad (20)$$

where

$$\tilde{\psi}_\tau(x, y) = \frac{(-1)^\tau}{2^{\tau-1}(\tau-1)!} \left(\frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2} \right)^\tau e^{-\frac{x^2+y^2}{2}}. \quad (21)$$

The integer parameter τ controls the number of vanishing moments and thus the shape of the wavelet. The filter response in the frequency domain broadens with decreasing τ . Our tests showed that values of $\tau > 2$ have to be avoided, because the filter selectivity becomes too low. Setting $\tau = 1$ has been found to be an appropriate value for our application. The corresponding wavelet is also known as the *Log* wavelet or Mexican hat wavelet, i.e. the Laplacian of a Gaussian. Its frequency response is given by:

$$\hat{\psi}(r) = r^2 e^{-\frac{r^2}{2}}. \quad (22)$$

For the directional separation of this isotropic wavelet, we shape it in angular direction in the frequency domain:

$$\hat{\psi}_{jk}(r, \theta) = \hat{\psi}_j(r) \cdot \eta_k(\theta). \quad (23)$$

The shaping function $\eta_k(\theta)$ used here is based on Schwarz functions, which are infinitely differentiable and of compact support

[7]. Most importantly, they can be normalized and combined to form the shaping function such that it satisfies

$$\sum_{k=0}^{K-1} |\eta_k(\theta - 2\pi k/K)|^2 = 1 \quad (24)$$

in accordance with Eq. (17). The number of filter orientations K is the parameter. The minimum number required by the analytic filter constraints, i.e. an angular support smaller than π , is three orientations. The human visual system emphasizes horizontal and vertical directions, so four orientations should be used as a practical minimum. To give additional weight to diagonal structures, eight orientations are preferred. One such analytic filter and the squared sum of all eight directional filters are shown in Figure 3. Although using even more filters might result in a better analysis of the local neighborhood, our experiments indicate that there is no apparent improvement when using more than eight orientations.




Fig. 3. One of eight analytic directional filters $\hat{\psi}_{jk}$ used for the contrast computation (left), and the isotropic sum of all eight (right) in the frequency domain.

Finally, the level j in the pyramidal decomposition has to be chosen. At any individual location the watermark may be considered a high-frequency distortion. As mentioned before, masking is strongest when masker and target have similar frequencies. Besides, contrast sensitivity is lowest for high frequencies. Therefore, we focus our efforts on this region, because it provides the greatest margin for hiding information in the image. In the filters presented above, the high-frequency information of the image is contained in the low levels. Figure 1 shows that level 0 best emphasizes high-frequency areas in the image, whereas the higher levels tend to smear the local contrast. The lowest level is therefore most suitable for the computation of contrast in our application.

C. Watermark Weights

As described by Eq. (3), the watermark is weighted with the function $\alpha(x, y)$. This function takes into account all the visual effects presented above. Once the isotropic contrast C_0^I from Eq. (17) and the corresponding visibility thresholds C_T from Eq. (18) have been computed, $\alpha(x, y)$ is given by:

$$\alpha(x, y) = C_T(C_0^I)(x, y) \cdot \phi_0 * I(x, y). \quad (25)$$

The local amplitude of the watermark at the threshold of visibility is thus determined by the multiplication of the contrast values with the corresponding low-pass filtered image ($\phi_0 * I$).

Figures 4 and 5 show examples for the images *lena* and *mandrill*, respectively. The dark and bright areas of the images in the left column represent the threshold and masking range of the masking model from Eq. (18), respectively. They illustrate the

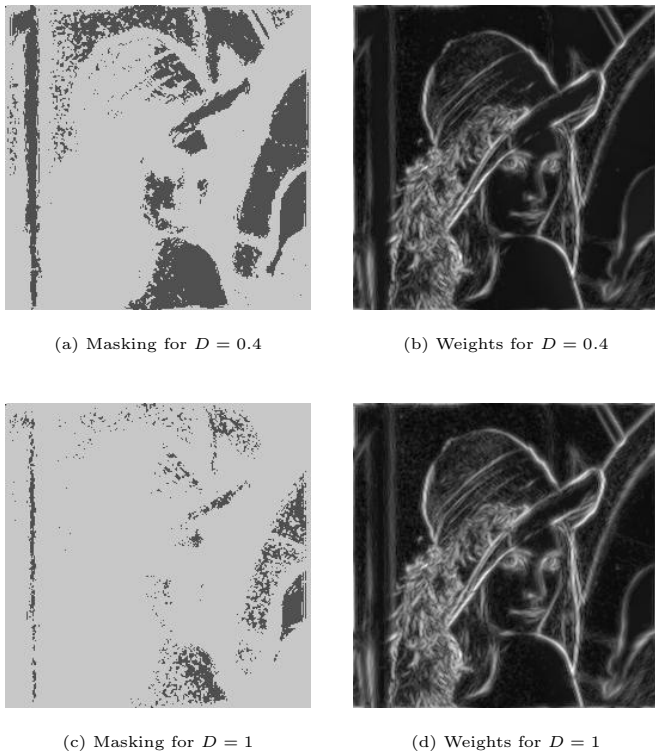


Fig. 4. Masking model and watermark weights of the *lena* image for two different density values. The dark and bright areas of the images in the left column represent the threshold and the masking range of the masking model, respectively. The right column shows the corresponding watermark weights.

assignment of texture and edge regions to the masking range, and the more uniform areas to the threshold range. One can clearly see that the watermark is weighted based on the masking phenomenon in most areas of the image. The right column of the same figures shows the resulting weighting functions $\alpha(x, y)$.

D. Wiener Filter

As we have seen in Section II-B, the Wiener filter needs an estimate of the local noise variance. Initially, an estimate of the original watermark weighting function $\alpha(x, y)$ was used to compute the variance estimate. However, our experiments indicated that this approach is not very efficient for our particular weighting function based on the masking model. This is due to the fact that under strong attacks, the amplitude of the watermark is heavily attenuated, and we have to detect a weak signal in additive noise. Moreover, it can be assumed that large watermark amplitudes are attenuated more than smaller amplitudes. It was found that the magnitude of the watermark after a strong attack is approximately proportional to the local luminance. From this we derived an efficient detector based on a linear corrector and a linear weighting function of the form $\tilde{\alpha}(x, y) = 1 + 4l(x, y)$, where $l(x, y)$ is the normalized local luminance. If the original weighting function were used, the variance of the watermark would be greatly overestimated, resulting in a notable decrease of the Wiener filter performance.

Based on these estimates, a Wiener filter of size 5×5 is employed in the detection process. This size exhibits the same performance as a 3×3 filter but has the added advantage of resisting to attacks with very small low-pass filters [10].

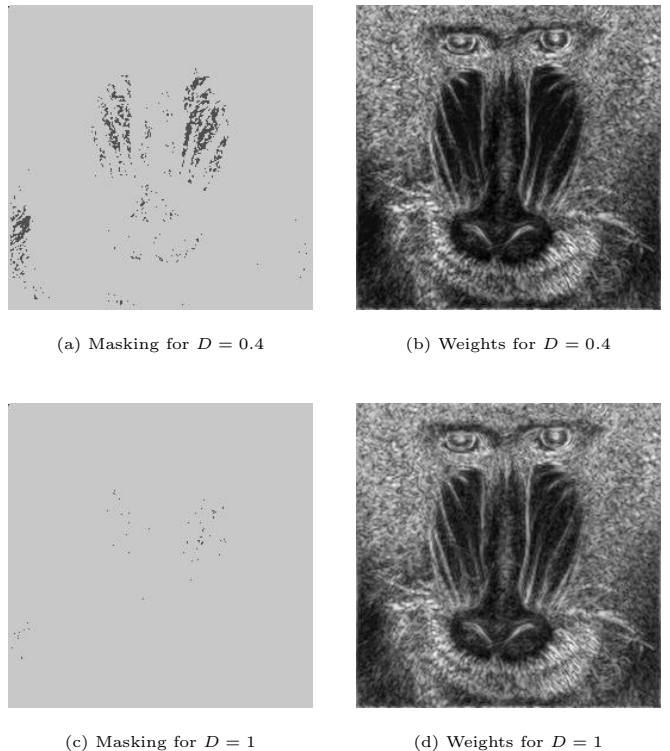


Fig. 5. Masking model and watermark weights of the *mandrill* image for two different density values. The dark and bright areas of the images in the left column represent the threshold and the masking range of the masking model, respectively. The right column shows the corresponding watermark weights.

V. SUBJECTIVE EXPERIMENTS

The masking model defined by Eq. (18) contains parameters ε , C_{T_0} and C_{M_0} . Since suitable experimental data on the masking of white noise do not exist in the scientific literature, we carried out a number of visual experiments to determine these parameters.

The setup for these experiments was as follows. The images *lena* and *mandrill* were displayed at a size of 256×256 pixels on a computer screen at a viewing distance of about 40 cm, thus extending over roughly 12 degrees of visual angle. Three subjects with normal or corrected-to-normal vision were used for the experiments. All tests were repeated for watermark densities D ranging from 0.1 to 1 in steps of 0.1. To allow implementations for all densities, the data were then approximated in a least-square sense. The specific functions for these fits were chosen independently based on each set of empirical data.

A. Luminance Masking

Embedding the watermark into the luminance channel is the most common approach. Luminance masking data from [5] suggest that $C_{M_0} \approx C_{T_0}$, which was confirmed by our experiments. Therefore only C_{T_0} and ε remain to be identified. These parameters were determined as follows. First C_{T_0} was measured by varying the noise amplitude uniformly across the entire image regardless of the image contrast. The results of this experiment are shown in Figure 6, where the detection threshold is plotted for each subject. It clearly demonstrates that the detection threshold increases for smaller densities.

After identifying the detection threshold, ε is varied to de-




Fig. 6. Detection threshold as a function of the density for luminance noise. Experimental data for three subjects and a least-square approximation are shown.

termine the slope of the masking function. For the different densities, the corresponding average detection thresholds from the previous test were used. This time, the subjects were asked to look for artifacts in texture areas and around edges. The test results are shown in Figure 7. It can be seen that ε decreases with increasing density.




Fig. 7. Exponent of the masking model as a function of the density for luminance noise. Experimental data for three subjects and a least-square approximation are shown.

B. Blue-Channel Masking

The human visual system is much less sensitive to blue colors than to others. This prompts the idea of watermarking the blue channel of a color image instead of the luminance signal, with the goal of embedding a watermark with higher energy at equivalent visual image quality. Nevertheless, applying the concept of masking to the blue channel requires empirical verification. It is well known that there is an interaction between luminance and color channels [21]. Computing the local contrast in the blue channel would not only cause problems with images containing no or little blue, e.g. computer generated images, but it would also neglect inter-channel masking. Therefore it is more suitable to use the luminance channel to compute the local contrast.

To evaluate the parameters for watermarking in the blue channel, the same approach as for the luminance case was investigated. First, the detection threshold for white noise was determined. The results of this experiment are shown in Fig-

ure 8. It is interesting to note that the detection threshold for blue noise is an order of magnitude larger than for luminance noise. This nicely illustrates the low sensitivity of the human visual system to blue stimuli.




Fig. 8. Detection threshold as a function of the density for blue-channel noise. Experimental data for three subjects and a least-square approximation are shown.

The second series of tests to determine ε showed that $C_{M_0} \gg C_{T_0}$, contrary to the luminance-only case. We found that C_{M_0} had approximately the same value as in the luminance case. This makes sense because the contrast computation is actually based on the luminance. The contrast computed from the luminance image appears to best represent the overall local activity in the image. This result is also consistent with experimental data from [21].

The results of the subjective tests for determining ε are shown in Figure 9. As in the luminance case, ε decreases with increasing noise density, but ε is significantly smaller for the masking of blue noise than for luminance masking.




Fig. 9. Exponent of the masking model as a function of the density for blue-channel noise. Experimental data for three subjects and a piecewise linear approximation are shown.

VI. RESULTS

A. Watermark Energy

We have seen that the two proposed schemes, luminance watermarking and blue-channel watermarking, have substantially different masking parameters. To compare the schemes, let us take a look at their respective noise energies and corresponding

distortions for the *lena* image. The normalized noise is computed as follows:

$$\sigma_N = \sqrt{\frac{1}{XY} \sum_x \sum_y |w(x, y)|^2}. \quad (26)$$

Figure 10 shows the normalized noise for the luminance and the blue-channel watermark as a function of the density. The curves clearly demonstrate that the blue-channel watermark contains substantially (on average 50 times) more energy than the luminance watermark. Furthermore, the energy remains nearly constant over the range of densities. This means that the watermark visibility is mainly a function of the watermark energy and is independent of the density.

Figure 11 shows the distortion in terms of dB for the two watermarking schemes, measured for each embedding channel. The larger distortion in the blue channel due to the higher watermark energy is evident. For the same visual quality, the distortion in the blue channel can be 16-18 dB higher than the distortion in the luminance channel. Like the watermark energy, the distortion remains approximately constant for the different embedding densities.




Fig. 10. Normalized watermark noise as a function of the density for luminance and blue-channel watermarking.




Fig. 11. Image distortion due to the watermarking process as a function of the density for luminance and blue-channel watermarking.

It should be noted that the proposed watermark weighting function was designed for embedding an invisible watermark at

detection threshold. However, in certain applications a reduction of visual image quality may be accepted in exchange for increased robustness. The proposed mask can easily be modified to incorporate such a variable strength watermark. The modification consists of a simple scaling of the mask, which may be considered as a vertical shift of the masking function. Intuitively this makes sense because it means that the detection threshold is increased uniformly, resulting in an overall increase of the watermark visibility.

B. Robustness

We have seen that at equivalent visual distortion, the blue-channel watermark contains substantially higher energy than the luminance watermark. We now investigate the implication of this difference by analyzing the performance of both schemes in terms of their robustness.

When the watermarked image is not manipulated or attacked, it is evident that blue-channel watermarking systems have a smaller detection error probability than luminance watermarking systems. This is also the case if the image undergoes geometrical transformations, is low-pass filtered, or subjected to additive noise. For such alterations it is obviously advantageous to use blue-channel watermarking. On the other hand, we have to investigate the effects of nonlinear attacks, such as lossy JPEG compression. The resistance to JPEG compression may be considered a comprehensive evaluation because the algorithm is adapted to the HVS. For example, JPEG compression distorts the blue channel more than others because the human visual system is less sensitive to it. Our blue-channel watermark is based on a similar premise. It should be noted that other nonlinear attacks may lead to different results, however.

To investigate the robustness of our watermarking system with respect to JPEG compression, tests were conducted with the images *lena* and *mandrill*, both of size 256×256 . The images were watermarked according to the luminance and the blue-channel schemes using embedding densities of $D = 0.5$ and $D = 1$. The length of the embedded signature was set to $N = 64$ bits. The value of the watermark was set to decimal 1234567890. In order to increase the statistical significance of the test results, the tests were repeated 50 times using a different key for each test. The JPEG quality setting was varied from 20% to 100% in steps of 5%.

The results for the images *lena* and *mandrill* are shown in Figures 12 and 13, respectively. The graphs show the detection error probability as a function of the JPEG quality setting for two watermark embedding densities. For the *lena* image we observe that the blue-channel watermark can accommodate lower quality settings (higher compression ratios) than the luminance watermark for both density settings. For the *mandrill* image, the behavior is different. The luminance watermark and blue-channel watermark at $D = 1$ and the luminance watermark at $D = 0.5$ feature about the same robustness, whereas the blue-channel watermark at $D = 0.5$ exhibits a slightly inferior performance. Complete detection failure, i.e. not detecting the watermark at all, occurs at approximately the same settings for both schemes. It is also interesting to note that higher robustness is not necessarily achieved with a higher watermark embedding density. This will be investigated in more detail in the next section.

It is clear that the introduced sparse spread-spectrum watermarking scheme is not resilient to geometrical alterations such as scaling, rotation and cropping. As mentioned in the introduction, the goal of the paper is the introduction of a new weighting function and not the design of an overall robust digital water-

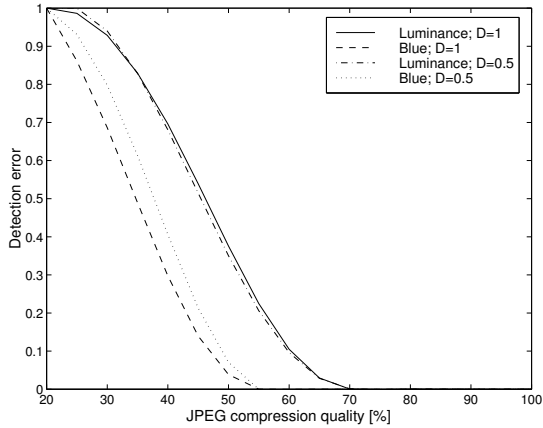


Fig. 12. Comparison of the robustness of luminance and blue-channel watermarking at two embedding densities for the *lena* image.

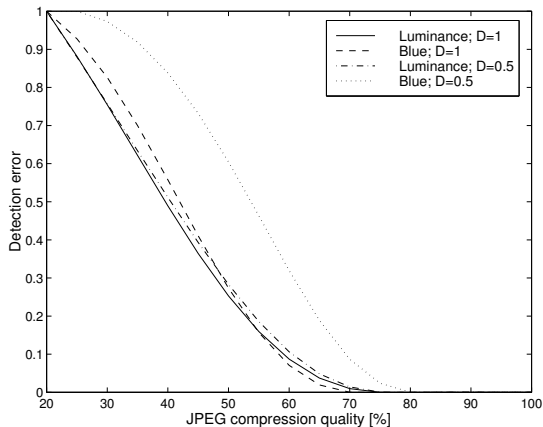


Fig. 13. Comparison of the robustness of luminance and blue-channel watermarking at two embedding densities for the *mandrill* image.

marking scheme. However, there are various possible extensions to the proposed scheme which can provide this functionality. These extensions are based on the inclusion of a reference watermark for spatial synchronization and on the concept of self-reference, i.e. multiple watermark embedding, for resilience to affine transformations [10].

C. Embedding Density

Most spread-spectrum based watermarking schemes employ a watermark embedding density of $D = 1$, i.e. they do not use a generalized scheme with variable density as proposed in this paper. Considering the results derived so far, we question this approach and ask if it may be better to use a lower density in exchange for a higher amplitude watermark.

In order to evaluate the impact of the density on the performance, various tests were conducted using the images *lena* and *mandrill*, both of size 256×256 . The experimental setup was the same as in the previous section. The tests were repeated for density values ranging from 0.1 to 1 in steps of 0.1.

The results for *lena* are shown in Figure 14. Comparing the curves for luminance and blue-channel schemes, we first observe the higher robustness of the blue-channel watermark for the entire range of density values. This confirms the results from the previous section. Furthermore, we observe that higher densities tend to result in higher robustness. However, the best performance is not achieved at the maximum density of $D = 1$. We





Fig. 14. Impact of the watermark embedding density on the robustness to lossy JPEG compression for the *lena* image. The two graphs show the detection error probability as a function of the embedding density and the JPEG quality setting for luminance (top) and blue-channel (bottom) watermarking.

may therefore deduce that the watermark interferes with itself at very high embedding density values, which results in a slight performance decrease. For the *lena* image, the highest overall robustness is achieved with an embedding density of $D \approx 0.8$.

The results for *mandrill* are shown in Figure 15. Again the general tendency of increased robustness for larger embedding densities becomes evident. For this image and the blue-channel watermark, the best performance is actually obtained for maximum density. Overall, however, densities close to 1 result in approximately the same performance. As before it should be noted that the results might be different for nonlinear attacks other than JPEG compression.

VII. CONCLUSIONS

We presented a novel weighting method for spread-spectrum watermarks. The approach is based on the human visual system and takes into account its masking properties in order to minimize the visual distortion and to increase the robustness of the watermark. The masking model necessitates a local isotropic contrast measure, which we compute from a pyramidal decomposition using directional analytic filters. The contrast measure and the masking model form the basis for the computation of a weighting function for the spread-spectrum watermark. Subjective experiments were carried out to find the optimal model parameters that determine the relationship between the density and the amplitude of the watermark. The resulting weighting function facilitates the insertion of a higher-energy watermark

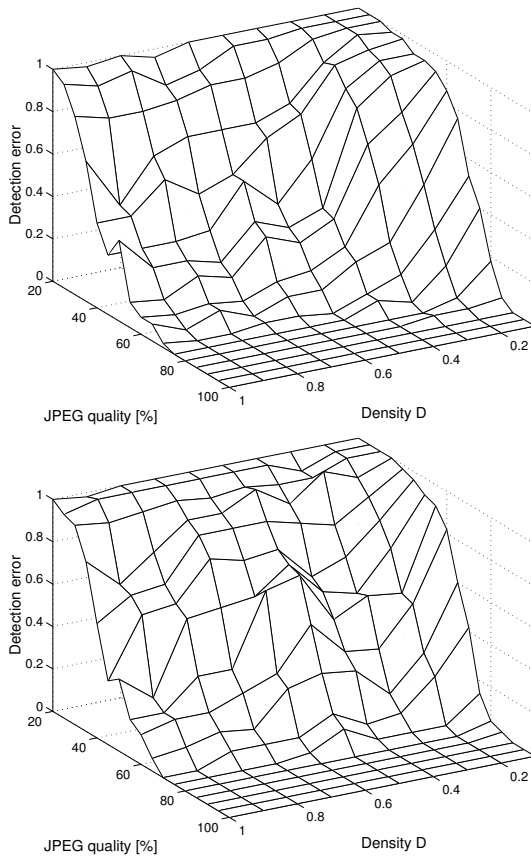


Fig. 15. Impact of the watermark embedding density on the robustness to lossy JPEG compression for the *mandrill* image. The two graphs show the detection error probability as a function of the embedding density and the JPEG quality setting for luminance (top) and blue-channel (bottom) watermarking.

while maintaining the same low visual distortion as other techniques. This leads to an increased robustness to various attacks.

The proposed method was used to watermark luminance images as well as the blue channel of color images. For equal image fidelity, the amplitude of the watermark must be adapted to the embedding density. In either case, the energy of the resulting watermark at the threshold of visibility remains approximately constant for all densities, i.e. the watermark energy is independent of the embedding density. However, the blue channel can accommodate a watermark of much higher energy than luminance.

There is not a single best watermark embedding channel for all situations. Depending on the nature of the attack and the image, one scheme may perform better than the other. This is particularly true for nonlinear attacks such as JPEG compression. If the robustness to geometrical attacks, additive noise or linear filtering is the primary concern, blue-channel watermarking is the method of choice.

Finally, increased embedding densities generally result in increased watermark robustness. However, due to watermark self-interference, the best performance may not be achieved at the maximum embedding density, but slightly below that.

REFERENCES

- [1] J.-P. Antoine, R. Murenzi, P. Vandergheynst: "Directional wavelets revisited: Cauchy wavelets and symmetry detection in patterns." *Appl. Comp. Harm. Anal.* **6**(3):314–345, 1999.
- [2] S. Daly: "The visible differences predictor: An algorithm for the assess-

ment of image fidelity." in *Digital Images and Human Vision*, ed. A. B. Watson, pp. 179–206, MIT Press, 1993.

- [3] J. G. Daugman: "Uncertainty relation for resolution in space, spatial frequency, and orientation optimized by two-dimensional visual cortical filters." *J. Opt. Soc. Am. A* **2**(7):1160–1169, 1985.
- [4] J. F. Delaigle, C. D. Vleeschouwer, B. Macq: "Watermarking algorithm based on a human visual model." *Signal Processing* **66**(3):319–335, 1998.
- [5] J. M. Foley: "Human luminance pattern-vision mechanisms: Masking experiments require a new model." *J. Opt. Soc. Am. A* **11**(6):1710–1719, 1994.
- [6] J. M. Foley, Y. Yang: "Forward pattern masking: Effects of spatial frequency and contrast." *J. Opt. Soc. Am. A* **8**(12):2026–2037, 1991.
- [7] J.-F. Gobbers, P. Vandergheynst: "Directional wavelet frames: Design and algorithms." *IEEE Trans. Image Processing* 2002, to appear.
- [8] F. Hartung, M. Kutter: "Multimedia watermarking techniques." *Proc. IEEE* **87**(7):1079–1107, 1999.
- [9] M. Kutter: "Watermarking resisting to translation, rotation, and scaling." in *Proc. SPIE*, vol. 3528, pp. 523–431, San Jose, CA, 1998.
- [10] M. Kutter: *Digital Image Watermarking: Hiding Information in Images*. Ph.D. thesis, École Polytechnique Fédérale de Lausanne, Switzerland, 1999.
- [11] M. Kutter, F. Jordan, F. Bossen: "Digital watermarking of color images using amplitude modulation." *J. Electronic Imaging* **7**(2):326–332, 1998.
- [12] J. S. Lim: *Two-Dimensional Signal and Image Processing*. Prentice Hall, 1990.
- [13] J. Lubin: "A visual discrimination model for imaging system design and evaluation." in *Vision Models for Target Detection and Recognition*, ed. E. Peli, pp. 245–283, World Scientific Publishing, 1995.
- [14] E. Peli: "Contrast in complex images." *J. Opt. Soc. Am. A* **7**(10):2032–2040, 1990.
- [15] E. Peli: "In search of a contrast metric: Matching the perceived contrast of Gabor patches at different phases and bandwidths." *Vision Res.* **37**(23):3217–3224, 1997.
- [16] R. L. Pickholz, D. L. Schilling, L. B. Milstein: "Theory of spread-spectrum communications – a tutorial." *IEEE Trans. Comm.* **30**(5):855–884, 1982.
- [17] C. I. Podilchuk, W. Zeng: "Image-adaptive watermarking using visual models." *IEEE J. Selected Areas in Comm.* **16**(4):525–539, 1998.
- [18] W. F. Schreiber: "Image processing for quality improvement." *Proc. IEEE* **66**(12):1640–1651, 1978.
- [19] J. R. Smith, B. O. Comiskey: "Modulation and information hiding in images." in *Lecture Notes in Computer Science: Information Hiding*, vol. 1174, pp. 207–226, Springer-Verlag, 1996.
- [20] E. M. Stein, G. Weiss: *Introduction to Fourier Analysis on Euclidean Spaces*. Princeton University Press, 1971.
- [21] E. Switkes, A. Bradley, K. K. De Valois: "Contrast dependence and mechanisms of masking interactions among chromatic and luminance gratings." *J. Opt. Soc. Am. A* **5**(7):1149–1162, 1988.
- [22] S. Winkler, P. Vandergheynst: "Computing isotropic local contrast from oriented pyramid decompositions." in *Proc. ICIP*, vol. 4, pp. 420–424, Kyoto, Japan, 1999.

Martin Kutter received his MS degree in electrical engineering in 1996 from the University of Rhode Island, Kingston, USA, and his Ph.D. degree in electrical engineering in 1999 from the Swiss Federal Institute of Technology (EPFL), Lausanne, Switzerland. For his Ph.D. thesis entitled "Digital image watermarking: hiding information in images", he received the best thesis award from the EPFL. Dr. Kutter has created the digital WatermarkingWorld, an Internet forum for people active in digital watermarking, and is the owner of the watermarking mailing list with over 700 subscribers worldwide. In 2001 he co-founded

AlpVision SARL, a digital watermarking company located in Switzerland. His research interests include digital watermarking of multimedia data, cryptography, data compression, image/video processing, and visual special effects.

Stefan Winkler received the M.Sc. degree in electrical engineering from the University of Technology in Vienna, Austria, in 1996, and the Ph.D. degree in electrical engineering from the Swiss Federal Institute of Technology in Lausanne, in 2000, for work on vision modeling and video quality assessment. In 2001 he co-founded Genimedia SA, Switzerland, a company focusing on perceptual quality of service for multimedia applications. His research interests include image/video processing, watermarking, human visual perception and quality metrics.