Wavelet-based Contrast Computation and Application to Digital Image Watermarking

Pierre Vandergheynst, Martin Kutter, Stefan Winkler
Signal Processing Laboratory
Swiss Federal Institute of Technology
CH-1015 Lausanne, Switzerland
Pierre.Vandergheynst@epfl.ch, Martin.Kutter@kutter.ch, Stefan.Winkler@epfl.ch

ABSTRACT

We introduce an isotropic measure of local contrast for natural images that is based on analytic filters and present the design of directional wavelet frames suitable for its computation. We show how this contrast measure can be used within a masking model to facilitate the insertion of a watermark in an image while minimizing visual distortion.

1. INTRODUCTION

Working with contrast instead of luminance can facilitate numerous image processing and analysis tasks. In section 2, we review the shortcomings of existing contrast definitions and introduce an isotropic measure of local contrast based on the combination of directional analytic filters. In section 3, we then present a filter design technique using a special dyadic wavelet transform that is tailored to the constraints imposed by our contrast definition. The usefulness of this isotropic contrast measure is demonstrated within the framework of a spread-spectrum watermarking application in section 4.

2. ISOTROPIC LOCAL CONTRAST

2.1. Contrast Definitions

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, a property known as *Weber's law*. Contrast is a measure of this relative variation of luminance. Mathematically, *Weber contrast* can be expressed as:

$$C^W = \frac{\Delta L}{L}. \tag{1}$$

This definition 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 of symmetrical deviations ranging from $L_{\min}$ to $L_{\max}$, which are also very popular in vision experiments, *Michelson contrast*[1] is generally used:

$$C^M = \frac{L_{\max} - L_{\min}}{L_{\max} + L_{\min}}. \tag{2}$$

While these two definitions are good predictors of perceived contrast for the above-mentioned classes of simple stimuli, they fail when stimuli become more complex and cover a wider frequency range, for example Gabor patches.[2] It is also evident that neither of these simple global definitions is appropriate for measuring contrast in natural images, because the brightest and darkest points would determine the contrast of the entire image, whereas actual human contrast perception varies with the *local* average luminance.

In order to address these issues, Peli[3] proposed a *local band-limited contrast*:

$$C_j^P(x,y) = \frac{\psi_j * I(x,y)}{\phi_j * I(x,y)}, \tag{3}$$

where $I(x,y)$ is the input image, ψ_j is a band-pass filter at level j of a filter bank, and ϕ_j is the corresponding low-pass filter. The normalization by the low-pass signal takes into account the local luminance variations. Modifications of this contrast definition have been used in a number of vision models[4,5] and are in good agreement with psychophysical experiments on Gabor patches.[2]

2.2. In-phase and Quadrature Mechanisms

Peli's local contrast as defined above measures contrast only as incremental or decremental changes from the local background, which is analogous to the symmetric (in-phase) responses of vision mechanisms. However, a complete description of contrast for complex stimuli has to include the anti-symmetric (quadrature) responses as well.[6,7]

The problem is demonstrated in Figure 1, which shows the contrast C^P computed with an isotropic band-pass filter for the *lena* image. It can be observed that C^P does not predict perceived contrast well, as it varies between positive and negative values of similar amplitude at the border between bright and dark regions and exhibits zero-crossings right where the perceived contrast is actually highest.

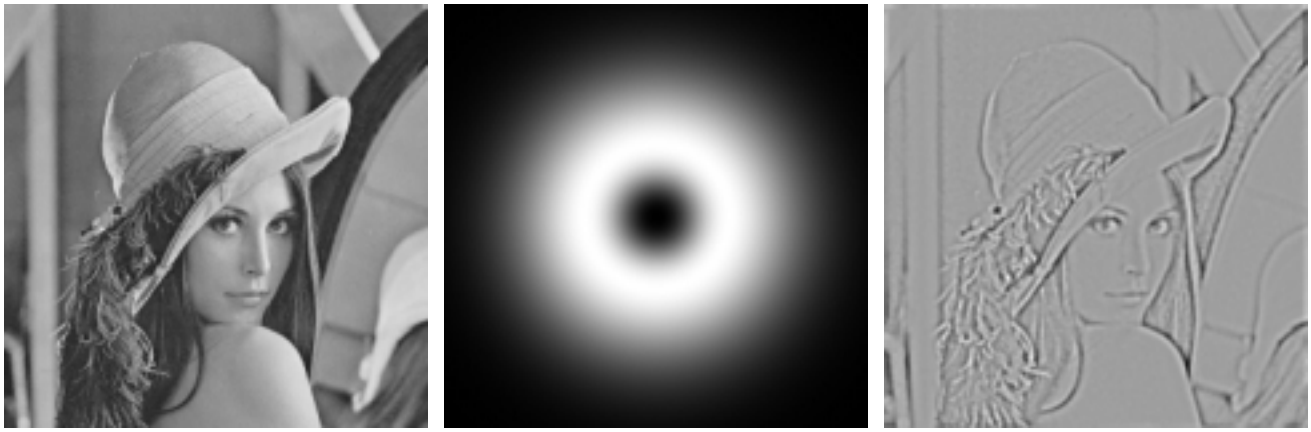

(a) *lena* image (b) Band-pass filter (c) Peli's contrast

Figure 1. Peli's local contrast from Eq. (3) computed for the *lena* image using an isotropic band-pass filter.

This behavior can be understood when C^P is computed for sinusoids with a constant C^M: The contrast computed using only a symmetric filter actually oscillates between $\pm C^M$ with the same frequency as the underlying sinusoid, which complicates establishing a correspondence between such a local contrast measure and data from psychophysical experiments.

These examples underline the need for taking into account both the in-phase and the quadrature component in order to be able to relate a generalized definition of contrast to the Michelson contrast of a sinusoid grating. Analytic filters represent an elegant way to achieve this: The magnitude of the analytic filter response, which is the sum of the energy responses of in-phase and quadrature components, exhibits the desired behavior, i.e. it gives a constant response to sinusoid gratings.

While the implementation of analytic filters in the one-dimensional case is straightforward, the design of general two-dimensional analytic filters is less obvious because of the difficulties involved in extending the Hilbert transform to two dimensions.[8] This problem is addressed in section 2.3 below.

Oriented measures of contrast can still be computed in two dimensions, though, because the Hilbert transform for filters whose angular support is smaller than π is well-defined. Such contrast measures can be used to implement a complete multi-channel representation of low-level vision, which has proven useful for modeling visual phenomena such as pattern masking. They have been employed in several vision models and their applications, e.g. in perceptual quality assessment of images and video,[9] or in image compression.[10]

2.3. Isotropic Local Contrast

We now present the construction of an *isotropic* contrast measure by combining analytic oriented filter responses.[11] It is the most natural measure of local contrast in an image. Isotropy is essential for applications where non-directional signals in an image are considered, e.g. spread-spectrum watermarking (see section 4).

The main problem in defining an isotropic contrast measure based on filtering operations is that if a flat response to a sinusoidal grating as with Michelson's definition is desired, 2-D analytic filters must be used. As already stressed in the previous section, the main difficulty in designing 2-D analytic filters is that there is no real equivalent to the Hilbert transform in two dimensions. Instead, there is a series of so-called *Riesz transforms*,[8] which are difficult to handle in practice, however.

In order to circumvent these problems, we present an approach using a class of non-separable filters that generalize the properties of analytic functions in 2-D. These filters are actually directional wavelets as defined by Antoine et al.,[12] which are square-integrable functions whose Fourier transform is strictly supported in a convex cone with the apex at the origin. It can be shown that these functions admit a holomorphic continuation in the domain $T = \mathbb{R}^2 + jV$, where V is the cone defining the support of the function. This is a genuine generalization of the Paley-Wiener theorem for analytic functions in one dimension. Furthermore, if we require that these filters have a flat response to sinusoidal stimuli, it suffices to impose that the opening of the cone V be less than π. This means that at least three such filters are required to cover all possible orientations uniformly, but otherwise there is no restriction on the number of filters. Using a technique described in section 3, such filters can be designed in a very simple and straightforward way; it is even possible to obtain dyadic oriented decompositions that can be implemented using a filter bank algorithm.

Working in polar coordinates (r, φ) in the Fourier domain, assume a directional wavelet $\hat{\Psi}(r, \varphi)$ satisfying the above requirements and

$$\sum_{k=1}^{K} |\hat{\Psi}(r, \varphi - 2\pi k/K)|^2 = |\hat{\psi}(r)|^2, \tag{4}$$

where K is the number of orientations, and $\hat{\psi}(r)$ is the Fourier transform of an isotropic dyadic wavelet, i.e.

$$\sum_{j=-\infty}^{+\infty} |\hat{\psi}(2^j r)|^2 = 1 \tag{5}$$

and

$$\sum_{j=-J}^{+\infty} |\hat{\psi}(2^j r)|^2 = |\hat{\phi}(2^J r)|^2, \tag{6}$$

where ϕ is the associated 2-D scaling function.[13]

Now it is possible to construct an isotropic contrast measure from the energy sum of directional filter responses:[11]

$$C_j^I(x, y) = \frac{\sqrt{2 \sum_k |\Psi_{jk} * I(x, y)|^2}}{\phi_j * I(x, y)}, \tag{7}$$

where Ψ_{jk} denotes the wavelet dilated by 2^{-j} and rotated by $2\pi k/K$. If the directional wavelet Ψ is in $L^1(\mathbb{R}^2) \cap L^2(\mathbb{R}^2)$, the convolution in the numerator of Eq. (7) is again a square-integrable function, and Eq. (4) shows that its L^2-norm is exactly what would have been obtained using the isotropic wavelet ψ. C_j^I is thus an orientation- and phase-independent quantity, but being defined by means of analytic filters, it behaves as prescribed with respect to sinusoidal gratings (i.e. $C_j^I(x, y) \equiv C^M$ in this case).

Examples for this isotropic contrast are shown in Figure 2. It can be seen that the contrast features obtained with C_j^I correspond very well to perceived contrast. The combination of the directional analytic filter responses produces a naturally meaningful phase-independent measure of isotropic contrast.

3. DIRECTIONAL WAVELET FRAMES

As stressed in Section 2.3, the computation of a robust isotropic contrast measure requires the use of a translation invariant multiresolution representation based on 2-D analytic filters. This can be achieved by designing a special Dyadic Wavelet Transform (DWT) using 2-D non-separable frames. The very weak design constraints of these frames permit the use of analytic wavelets, and the isotropy condition (4) can easily be fulfilled.

(a) Level 0 (b) Level 1 (c) Level 2

Figure 2. Isotropic contrast of the *lena* image as described by Eq. (7) at three different levels.

3.1. 2-D Wavelet Frames

In the following we will always make use of a set of K oriented bidimensional wavelets $\psi^k(\vec{x})$, $1 \leq k \leq K$, such that their frequency support roughly spans a corona. If we write scaled and translated wavelets:

$$\psi^k_{2^j,\vec{b}}(\vec{x}) = \frac{1}{2^{2j}}\psi^k\Big(2^{-j}\big(\vec{x}-\vec{b}\big)\Big)\,, \tag{8}$$

the Dyadic Wavelet Transform (DWT) of $f \in L^2\big(\mathbb{R}^2, d^2\vec{x}\big)$ is given by

$$\mathcal{W}^k_f\big(2^j,\vec{b}\big) = \langle \psi^k_{2^j,\vec{x}}, f\rangle = \big(\tilde{\psi}^k * f\big)(\vec{x}) \tag{9}$$

with $\tilde{\psi}(\vec{x}) = \overline{\psi(-\vec{x})}$. The DWT is a complete and stable representation[14] provided there exist two positive finite constants A and B such that

$$A \leq \sum_{k=1}^{K}\sum_{j\in\mathbb{Z}} |\hat{\psi}^k(2^j\vec{\omega})|^2 \leq B \quad \forall\vec{\omega}\in\mathbb{R}^2\setminus\vec{0}\,. \tag{10}$$

When (10) holds, one may build reconstruction or dual wavelets

$$\hat{\chi}^k(\vec{\omega}) = \frac{\hat{\psi}^k(\vec{\omega})}{\sqrt{\sum_k\sum_j |\hat{\psi}^k(2^j\vec{\omega})|^2}}, \tag{11}$$

for which we have

$$\sum_{j\in\mathbb{Z}}\sum_{k=1}^{K} \hat{\chi}^k(2^j\vec{\omega})\overline{\hat{\psi}^k(2^j\vec{\omega})} = 1\,. \tag{12}$$

Finally the following reconstruction formula holds in L^2 sense:

$$f(\vec{x}) = \sum_{j\in\mathbb{Z}} 2^{-2j}\sum_{k=1}^{K}\Big(\mathcal{W}^k_f\big(2^j,\cdot\big) * \chi^k_{2^j}\Big)(\vec{x})\,. \tag{13}$$

Replacing (10) by the weaker condition

$$\sum_{k=1}^{K}\sum_{j\in\mathbb{Z}} \hat{\psi}^k(2^j\vec{\omega}) = 1 \quad \forall\vec{\omega}\in\mathbb{R}^2\setminus\vec{0}\,, \tag{14}$$

a simpler reconstruction formula holds:

$$f(\vec{x}) = \sum_{j\in\mathbb{Z}} \sum_{k=1}^{K} 2^{-j} \mathcal{W}_f^k\big(2^j, \vec{x}\big) \ . \tag{15}$$

One of the main advantages of using the DWT lies in the possibility of working with translation invariant frames. This is a very important requirement in many areas of computer vision where one wants to study primitives that do not depend on positions in the image. In particular, the estimation of luminance contrast discussed in the previous section should be translation invariant, in agreement with physical reality. Other useful properties that can be easily achieved with dyadic frames are isotropy or symmetry, fast exponential decay, prescribed smoothness etc.

3.2. Custom Design of Dyadic Frames

Even though equations (10,11) are weaker than usual design conditions for an orthogonal basis, it is still somewhat tricky to efficiently tailor dyadic frames to match specific needs. In this section we briefly review a technique based on discretizing a continuous wavelet transform.[15,16]

We start from the reconstruction formula of the 2-D Continuous Wavelet Transform (CWT):[17]

$$f(\vec{x}) = \int_0^{+\infty} a^{-1} da \int_{\mathbb{R}^2} d^2\vec{b}\ \mathcal{W}_f(a,\vec{b}) \psi_{\vec{b},a}(\vec{x}) \ , \tag{16}$$

where $\mathcal{W}_f(a,\vec{b})$ is the wavelet transform of the signal $f \in L^2\big(\mathbb{R}^2, d^2\vec{x}\big)$,

$$\mathcal{W}_f(a,\vec{b}) = C_\psi^{-1} a^{-2} \int_{\mathbb{R}^2} d^2\vec{x}\ \overline{\psi\big(a^{-1}(\vec{x}-\vec{b})\big)} f(\vec{x}), \tag{17}$$

and C_ψ is a constant,

$$C_\psi = 2\pi \int_{\mathbb{R}^2} d^2\vec{\omega}\ \frac{|\hat{\psi}(\vec{\omega})|^2}{\|\vec{\omega}\|^2} < +\infty \tag{18}$$

that we will normalize to one in the following. These equations hold for an isotropic wavelet, but more general functions can be considered provided one adds rotations in (16) and (17). Defining the *infinitesimal detail*:

$$d_a(\vec{x}) = \int_{\mathbb{R}^2} d^2\vec{b}\ \mathcal{W}_f(a,\vec{b}) \psi_{\vec{b},a}(\vec{x}) \, , \tag{19}$$

equation (16) can be rewritten

$$f(\vec{x}) = \int_0^{+\infty} \frac{da}{a}\, d_a(\vec{x}) \ . \tag{20}$$

In order to obtain a discrete reconstruction formula, we want to replace the integral over scales in this equation by a discrete sum. Truncating this expansion to a given resolution, say a_0, we can synthesize a low-pass approximation of the signal:

$$f_{a_0}(\vec{x}) = \int_{a_0}^{+\infty} \frac{da}{a}\ d_a(\vec{x}) \, . \tag{21}$$

Taking the Fourier on both sides of this equation suggests to introduce the following Fourier multiplier:

$$|\hat{\phi}(\vec{\omega})|^2 = \int_1^{+\infty} \frac{da}{a}\ |\hat{\psi}(a\vec{\omega})|^2 \, . \tag{22}$$

This definition implies

$$\lim_{\|\vec{\omega}\|\to\infty} |\hat{\phi}(\vec{\omega})|^2 = 0, \tag{23}$$

which shows that ϕ is a smoothing function. Now starting from scale $a = 2^J$, we can refine up to an arbitrary resolution by adding up details. For this purpose, we introduce slices of details:

$$D_j(\vec{x}) = \int_{2^j}^{2^{j+1}} \frac{da}{a}\ d_a(\vec{x}) \, . \tag{24}$$

Computing the Fourier transform on both sides of this equation gives

$$\hat{D}_j(\vec{\omega}) = \hat{f}(\vec{\omega}) \int_{2^j}^{2^{j+1}} \frac{da}{a} \, |\hat{\psi}(a\vec{\omega})|^2, \tag{25}$$

which allows us to compute dyadic slices of wavelet coefficients with the following multiplier called *integrated wavelet packet*:

$$|\hat{\Gamma}(\vec{\omega})|^2 = \int_{\frac{1}{2}}^{1} \frac{da}{a} \, |\hat{\psi}(a\vec{\omega})|^2 \, . \tag{26}$$

Putting equations (26) and (22) together results in a dyadic decomposition:

$$f = \int_{\mathbb{R}^2} d^2\vec{b} \, \langle \phi_{(2^J,\vec{b})} | f \rangle \phi_{(2^J,\vec{b})} \, + \, \sum_{j=-\infty}^{J} \int_{\mathbb{R}^2} d^2\vec{b} \, \langle \Gamma_{(2^J,\vec{b})} | f \rangle \Gamma_{(2^J,\vec{b})} \, . \tag{27}$$

Obviously, Eq. (27) has to be understood in the strong sense in $L^2\big(\mathbb{R}^2, d^2\vec{x}\big)$.[15]

The basic element of this construction is the underlying continuous wavelet ψ. The integrated wavelet Γ and scaling function ϕ inherit its smoothness and localization properties. Due to the freedom we have in designing ψ * this gives a simple yet powerful way of deriving a translation invariant dyadic frame. Since this construction mainly works in the Fourier domain, it is very easy to add directional sensitivity by multiplying all Fourier transforms with a suitable angular window η. This leads to a polar decomposition of the form

$$\hat{\Psi}(r, \varphi) = \hat{\Gamma}(r)\eta(\varphi) \, . \tag{28}$$

In order to satisfy condition (4), it suffices to impose

$$\sum_{k=0}^{K-1} \Big| \eta\big(\varphi - \frac{2\pi k}{K}\big) \Big|^2 = 1. \tag{29}$$

Figure 3 shows an example of such a construction for $K = 5$ orientations. Such filters are ideally suited for the computation of isotropic contrast according to Eq. (7).

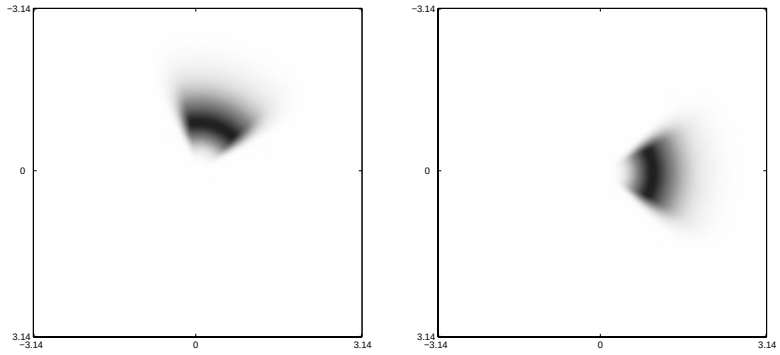

Figure 3. Fourier transform of a dyadic oriented wavelet with angular resolution of $\pi/5$ at two different orientations.

The main drawback of the technique described above is perhaps the lack of fast algorithms. In particular one would appreciate the existence of a pyramidal algorithm,[14] which is not guaranteed here because integrated wavelets and scaling functions are not necessarily related by a two-scale equation. On the other hand it has been demonstrated[16,18] that one can find approximate quadrature filters that achieve a fast implementation of the DWT with very accurate results. Once again, the advantage here is that it leaves us free to design our own dyadic frame.

* It only has to satisfy the admissibility condition (18).

4. APPLICATION TO IMAGE WATERMARKING

Digital watermarking is the art of invisible and robust data hiding in digital media. Despite the fact that we will focus on digital images in this paper, the reader should be aware that digital watermarking is not limited to this kind of media. For an in-depth overview of digital watermarking, the reader is referred to Hartung and Kutter.[19]

Watermarking describes the process of embedding information, usually in the form of a binary string, into an image by slightly manipulating the visual information. Embedding may take place in various domains, e.g. the spatial domain, the frequency domain, or the DCT domain. Robustness refers to the ability of the embedded watermark to survive image alterations, which can include geometrical transformations as well as the addition of noise, for example due to lossy compression or a print-scan procedure. The most important requirement, however, is watermark imperceptibility. The insertion of the watermark in the image has to be performed in such a way that the resulting modifications are not visible to the human eye. This is a very delicate and complex issue because it depends on a variety of interacting factors.

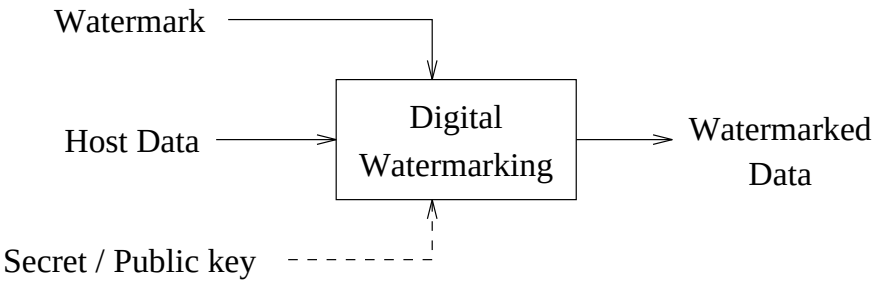

Figure 4. Generic digital watermarking scheme.

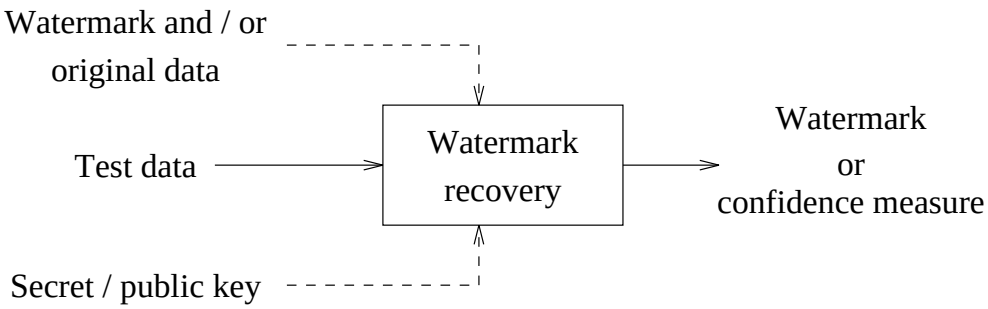

Figure 5. Generic watermark recovery scheme.

Figures 4 and 5 illustrate the basic concepts of watermark embedding and watermark recovery. Figure 4 shows the generic watermarking scheme for the embedding process. The inputs are the watermark, the *host data*, and an optional public or secret key. The host data may be uncompressed or compressed, depending on the application; most proposed methods work on uncompressed data, however. The watermark can be of any nature such as a number, text, or an image. The secret or public key is used to protect the watermark if it is not to be read by unauthorized parties. The output of the watermarking process is the modified, i.e. watermarked, data. The generic watermark recovery process is depicted in Figure 5. Inputs are the watermarked data, the secret or public key, and sometimes also the original data and/or the original watermark. The output of the recovery process is either the detected watermark or some kind of confidence measure indicating how likely it is for the given watermark to be present in the data.

From a mathematical point of view, digital watermarking is a constraint optimization problem. The goal is to maximize the watermark energy under the minimum visibility constraint. This optimization is non-trivial since it requires an accurate description of the human visual system in order to evaluate the visibility of the watermark. Although several vision models have been proposed in the past, most of them are very complex and often very specific to a given application.

In order to model the visibility of the watermark in the host image, there are mainly two effects that need to be taken into account, namely contrast sensitivity and masking. Contrast sensitivity describes the response of the human visual system to contrast as a function of the nature of the stimulus. Masking describes the phenomenon in which one signal, the masker, is capable of "hiding" a second signal, the target. In other words, the target visibility depends on the presence of a masker.

It is possible to combine contrast sensitivity and masking in the so-called *transducer model*. It is used to describe the relation between the masker contrast and the target contrast at detection threshold. Figure 6 shows the transducer function. On the horizontal axis we have the logarithm of the masker contrast C_M, and on the vertical axis we have the logarithm of the target contrast C_T. The 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 mathematical description of this model is given by:

$$C_T(C_M) = \begin{cases} C_{T_0} & \text{if } C_M < C_{M_0}, \\ C_{T_0}\,(C_M/C_{M_0})^{\varepsilon} & \text{otherwise.} \end{cases} \tag{30}$$

The model comprises three parameters, ε, C_{T_0} and C_{M_0}, which specify the size of the threshold and the masking range as well as the slope of the transducer function. They have to be determined by means of subjective experiments (see below).

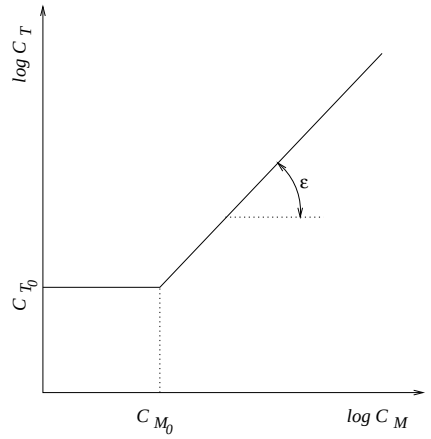

Figure 6. Transducer function describing the relationship between the masker contrast and the target contrast at detection threshold.

For evaluation purposes we apply this masking model to a watermarking scheme based on spatial spread-spectrum modulation as proposed by Kutter.[20] Each bit to be embedded in the image is represented by a two-dimensional pseudo-random pattern. The statistics of the pattern are bimodal with equal probabilities for -1 an 1. The random patterns of all bits are superimposed as follows:

$$w(x,y) = \alpha(x,y) \sum_i p_i(x,y), \tag{31}$$

where $p_i(x,y)$ are the pseudo-random modulation functions for bit i, $\alpha(x,y)$ is the watermark weighting function, and $w(x,y)$ is the resulting watermark which is added to the image. In this watermarking scheme the pseudo-random patterns p_i are sparse, which means that the superposition of all patterns does not necessarily modify all pixels in the image. To quantify the sparseness, we introduce the density D of the watermark, which is given by the modified number of pixels divided by the total number of pixels in the image.

To compute the watermark weighting function $\alpha(x,y)$, we make use of the introduced masking model and the local isotropic contrast measure presented in section 2.3. For computing the local contrast according to Eq. (7), we use directional wavelet frames as described in section 3 based on the *PLog* wavelet.[21] The minimum number of orientations required by the analytic filter constraint, i.e. an angular support smaller than π, is three. 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. Although using even more filters could result in a better analysis of the local neighborhood, our experiments indicate that there is no apparent improvement when using more than eight orientations. Also, we have to choose the frequency band of the pyramidal decomposition. At an individual location the watermark may be considered a high-frequency distortion. Furthermore, masking is strongest when masker and target have similar frequencies. Figure 2 shows that level 0 best emphasizes high-frequency areas in the image, whereas the higher levels tend to smear the local contrast. Therefore, the lowest level is most suitable for the computation of contrast in our application.

The watermark weighting function α is now computed as follows:

$$\alpha(x,y) = C_T(C_0^I)(x,y) \cdot \phi_0 * I(x,y), \tag{32}$$

where C_0^I is the local isotropic contrast of the masker image at level 0, C_T is the corresponding target contrast threshold as given by the transducer model, and ϕ_0 is a low-pass filter. The local amplitude of the watermark at the threshold of visibility is thus determined by the multiplication of the isotropic contrast values with the corresponding low-pass filtered image.

Next we need to determine the parameters of the transducer function, i.e. C_{T_0}, C_{M_0} and ε, through subjective tests. The experimental setup was based on displaying a color image of size 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 participated in the experiments. All tests were repeated for different watermark densities D ranging from 0.1 to 1 in steps of 0.1. Since we know from other luminance masking experiments[22] that $C_{M_0} \approx C_{T_0}$ (this was confirmed by our experiments as well), we are left with identifying C_{T_0} and ε.

These parameters are determined as follows. First we measure C_{T_0} by varying the noise amplitude uniformly across the entire image regardless of image contrast. The results of this experiment are shown in Figure 7(a), where the detection threshold is plotted for each subject. It clearly demonstrates that the detection threshold increases for smaller densities. To allow implementations for all densities, the data are approximated with an exponential function in a least-square sense.

After identifying the detection threshold, we vary ε to determine 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(b). The exponent ε is within the range typically found in masking experiments and decreases with increasing density.

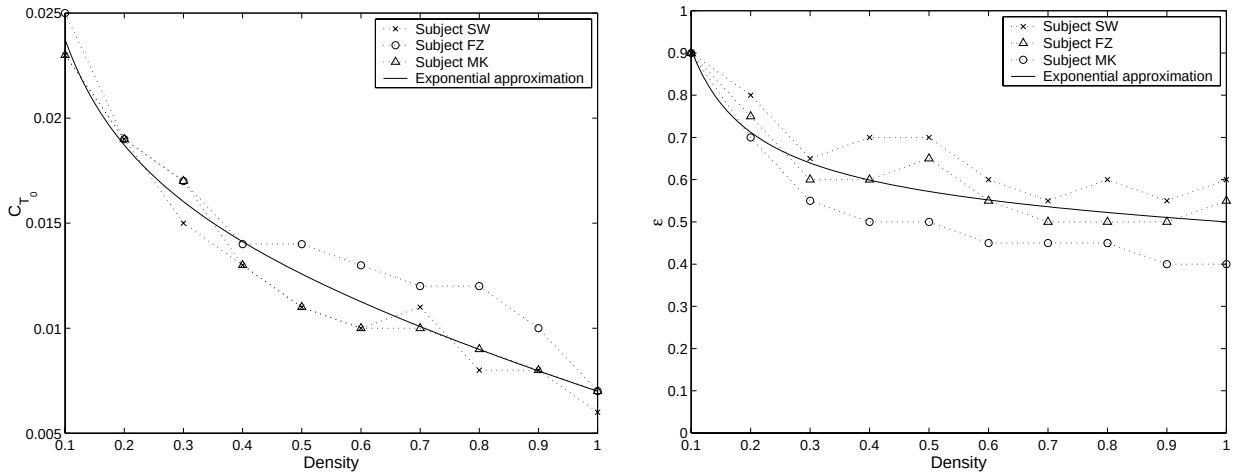

(a) Detection threshold

(b) Exponent of the masking model

Figure 7. Detection threshold C_{T_0} and exponent ε of the transducer model as a function of the density for luminance noise. Experimental data for three subjects and a least-square approximation are shown.

Using these parameters, the weighting mask $\alpha(x, y)$ can now be computed. Figure 8 shows weighting masks for the *lena* image at watermark densities of 0.4 and 1, respectively. For illustrative purposes the figures to the right visualize the segmentation into threshold and masking ranges. The dark areas correspond to regions where only contrast sensitivity is exploited, and the bright areas show regions where the masking effect is dominant.

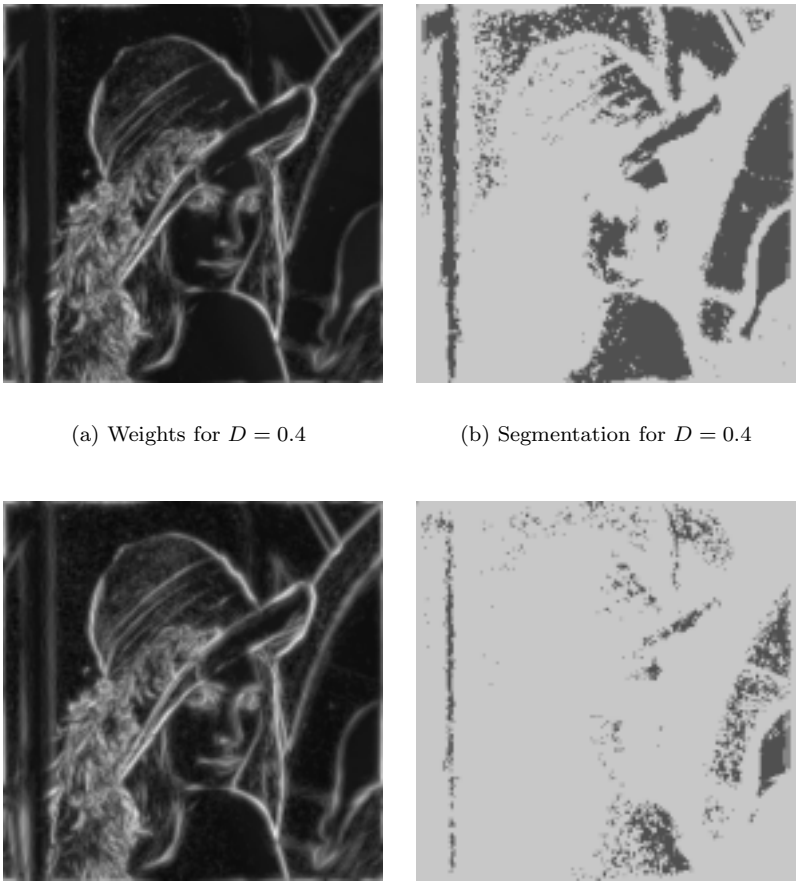

(a) Weights for $D = 0.4$

(b) Segmentation for $D = 0.4$

(c) Weights for $D = 1$

(d) Segmentation for $D = 1$

Figure 8. Watermark weighting function of the *lena* image for two different density values (left column). The segmented images (right column) illustrate the threshold and masking ranges of the transducer model, represented by dark and bright areas, respectively.

In comparison to other watermarking schemes, this weighting mask based on the simple masking model presented above facilitates the insertion of a watermark with higher energy while preserving the visual quality of the image, leading to a watermark that is more robust. It has also been applied successfully to watermarking the blue channel of color images.[21]

5. CONCLUSIONS

We presented a technique for measuring isotropic local contrast in natural images. It has many potential applications and can be used inside vision models, as demonstrated in section 4. The wavelet formulation of our contrast definition allows for covariance under translations and dyadic scalings. Finally, its subband or multiresolution nature permits more tuning, such as frequency weighting or luminance adaptation, which will be the subject of future work.

REFERENCES

1. A. A. Michelson, *Studies in Optics*, University of Chicago Press, 1927.
2. E. Peli, "In search of a contrast metric: Matching the perceived contrast of Gabor patches at different phases and bandwidths," *Vision Res.* **37**(23), pp. 3217–3224, 1997.
3. E. Peli, "Contrast in complex images," *J. Opt. Soc. Am. A* **7**, pp. 2032–2040, Oct. 1990.
4. S. Daly, "The visible differences predictor: An algorithm for the assessment of image fidelity," in *Digital Images and Human Vision*, A. B. Watson, ed., pp. 179–206, MIT Press, 1993.
5. J. Lubin, "A visual discrimination model for imaging system design and evaluation," in *Vision Models for Target Detection and Recognition*, E. Peli, ed., pp. 245–283, World Scientific Publishing, 1995.
6. 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**, pp. 1160–1169, July 1985.
7. C. F. Stromeyer, III. and S. Klein, "Evidence against narrow-band spatial frequency channels in human vision: The detectability of frequency modulated gratings," *Vision Res.* **15**, pp. 899–910, 1975.
8. E. M. Stein and G. Weiss, *Introduction to Fourier Analysis on Euclidean Spaces*, Princeton University Press, 1971.
9. S. Winkler, "Issues in vision modeling for perceptual video quality assessment," *Signal Processing* **78**, pp. 231–252, Oct. 1999.
10. P. Vandergheynst and Ö. N. Gerek, "Nonlinear pyramidal image decomposition based on local contrast parameters," in *Proc. Nonlinear Signal and Image Processing Workshop*, vol. 2, pp. 770–773, (Antalya, Turkey), June 20–23 1999.
11. S. Winkler and P. Vandergheynst, "Computing isotropic local contrast from oriented pyramid decompositions," in *Proc. ICIP*, vol. 4, pp. 420–424, (Kyoto, Japan), Oct. 25–28, 1999.
12. J.-P. Antoine, R. Murenzi, and P. Vandergheynst, "Directional wavelets revisited: Cauchy wavelets and symmetry detection in patterns," *Appl. Comp. Harm. Anal.* **6**, pp. 314–345, May 1999.
13. S. Mallat and S. Zhong, "Characterization of signals from multiscale edges," *IEEE Trans. PAMI* **14**, pp. 710–732, July 1992.
14. S. Mallat, *A Wavelet Tour of Signal Processing*, Academic Press, 1998.
15. M. Duval-Destin, M. A. Muschietti, and B. Torresani, "Continuous wavelet decompositions: Multiresolution and contrast analysis," *SIAM J. Math. Anal.* **24**, pp. 739–755, May 1993.
16. J.-F. Gobbers and P. Vandergheynst, "Directional wavelet frames: Design and algorithms," *IEEE Trans. Image Processing* , 1999. Submitted paper, preprint UCL-IPT-98-17.
17. J.-P. Antoine, P. Carrette, R. Murenzi, and B. Piette, "Image analysis with two-dimensional continuous wavelet transform," *Signal Processing* **31**, pp. 241–272, 1993.
18. M. A. Muschietti and B. Torrésani, "Pyramidal algorithms for Littlewood-Paley decompositions," *SIAM J. Math. Anal.* **26**, pp. 925–943, July 1995.
19. F. Hartung and M. Kutter, "Multimedia watermarking techniques," *Proc. IEEE* **87**, pp. 1079–1107, July 1999.
20. M. Kutter, *Digital Image Watermarking: Hiding Information in Images.* PhD thesis, École Polytechnique Fédérale de Lausanne, Switzerland, 1999.
21. M. Kutter and S. Winkler, "Spread-spectrum watermarking using the human visual system," *IEEE Trans. Image Processing* , 2000. Submitted paper.
22. J. M. Foley, "Human luminance pattern-vision mechanisms: Masking experiments require a new model," *J. Opt. Soc. Am. A* **11**, pp. 1710–1719, June 1994.