



Self-Invertible 2D Log-Gabor Wavelets

SYLVAIN FISCHER

Instituto de Optica, CSIC, Serrano 121, 28006 Madrid, Spain;

INCM, UMR6193, CNRS & Aix-Marseille University, 31 ch. Aiguier, 13402 Marseille Cedex 20, France

FILIP ŠROUBEK

Instituto de Optica, CSIC, Serrano 121, 28006 Madrid, Spain;

Academy of Sciences, Pod vodárenskou věží 4 Prague, Czech Republic

LAURENT PERRINET

INCM, UMR6193, CNRS & Aix-Marseille University, 31 ch. Aiguier, 13402 Marseille Cedex 20, France

RAFAEL REDONDO AND GABRIEL CRISTÓBAL

Instituto de Optica, CSIC, Serrano 121, 28006 Madrid, Spain

Received April 27, 2006; Accepted December 13, 2006

First online version published in January, 2007

Abstract. Orthogonal and biorthogonal wavelets became very popular image processing tools but exhibit major drawbacks, namely a poor resolution in orientation and the lack of translation invariance due to aliasing between subbands. Alternative multiresolution transforms which specifically solve these drawbacks have been proposed. These transforms are generally overcomplete and consequently offer large degrees of freedom in their design. At the same time their optimization gets a challenging task. We propose here the construction of log-Gabor wavelet transforms which allow exact reconstruction and strengthen the excellent mathematical properties of the Gabor filters. Two major improvements on the previous Gabor wavelet schemes are proposed: first the highest frequency bands are covered by narrowly localized oriented filters. Secondly, the set of filters cover uniformly the Fourier domain including the highest and lowest frequencies and thus exact reconstruction is achieved using the same filters in both the direct and the inverse transforms (which means that the transform is self-invertible). The present transform not only achieves important mathematical properties, it also follows as much as possible the knowledge on the receptive field properties of the simple cells of the Primary Visual Cortex (V1) and on the statistics of natural images. Compared to the state of the art, the log-Gabor wavelets show excellent ability to segregate the image information (e.g. the contrast edges) from spatially incoherent Gaussian noise by hard thresholding, and then to represent image features through a reduced set of large magnitude coefficients. Such characteristics make the transform a promising tool for processing natural images.

Keywords: wavelet transforms, log-Gabor filters, oriented high-pass filters, image denoising, visual system,

1. Introduction

After the development and rapid successes of the wavelet transforms in image processing, alternative multiresolutions have been proposed mainly for a better resolution in orientation and for avoiding aliasing effects. Last fifteen years oriented band-pass multiresolution transforms have arisen with increased importance thanks to the development of steerable pyramids (Simoncelli et al.,

1992), Gabor multiresolutions (Gross and Koch, 1995; Nestares et al., 1998; Ro et al., 2001), complex-valued wavelets (Gross and Koch, 1995; Kingsbury, 2001; Portilla and Simoncelli, 2000), curvelets (Starck et al., 2002) and contourlets (Do and Vetterli, 2005), to name a few.

A Gabor function is a Gaussian multiplied by a complex exponential. Thus, in the Fourier domain, it is a

Gaussian shifted from the origin. Gabor functions gather a number of interesting mathematical properties: first they have a smooth and infinitely differentiable shape. Second, their modulus are monomodal, i.e. they have no side lobes neither in space nor in the Fourier domain. Third, they are optimally joint-localized in space and frequency (Gabor, 1946). For such good properties they have been proposed as ideal functions for signal processing.

Similarly, 2D Gabor functions are highly jointly localized in position, orientation and spatial frequency. Neuroscience studies have shown that the receptive fields of simple cells of the Primary Visual Cortex (V1) of primates can be modelled by Gabor functions (Daugman, 1985; Marcelja, 1980). Considering natural vision as optimized by the millions of years of evolutionary pressure, the recruitment of Gabor functions by V1 can be understood as an additional argument in favor of their adequacy for image processing.

Nevertheless Gabor functions present some important drawbacks. First, it is not possible to build a complete orthogonal basis of Gabor functions, therefore non-orthogonal bases have to be employed. Non-orthogonality implies that exact reconstruction using the same filters for analysis and synthesis will not be possible unless an overcomplete basis is considered. Secondly, Gabors are bandpass filters, they are consequently inadequate for covering the lowest and highest frequencies. Thirdly, it is particularly difficult to cover up the mid frequencies with sufficient uniformity.

Gabor multiresolutions have been successfully used for image analysis and applications where exact reconstruction is not required, such as texture analysis (Clausi and Jernigan, 2000; Ro et al., 2001), texture synthesis (Portilla et al., 1996), edge/contour extraction (Heitger et al., 1998; Kovese, 2003; Grigorescu et al., 2003), or object recognition (Pöttsch et al., 1996; Krüger, 2001). And, even without exact reconstruction they have been shown useful for image restoration applications (Cristóbal and Navarro, 1994; Kovese, 1999; Mingolla et al., 1999; Christiansen, 2002). In parallel, different methods for reconstruction improvement have been proposed to recover the highest frequencies (Nestares et al., 1998), to avoid excessive low-pass overlapping (Field, 1987), to improve the reconstruction (Daugman, 1988; Lee, 1996) or to cover more uniformly the Fourier domain (Gross and Koch, 1995).

One of the important applications where oriented multiresolution schemes appeared to be particularly efficient is image denoising. In this field (bi-)orthogonal wavelets are themselves one of the most popular and optimal tools (Donoho, 1995). Nevertheless it has been shown that undecimated wavelets provide better results than the critically sampled wavelets (Coifman and Donoho, 1995; Chang et al., 2000; Li and Orchard, 2000). Latter works have shown the interest of redesigning the Fourier

domain tiling, particularly for increasing the number of orientations (Starck et al., 2002; Do and Vetterli, 2005; Portilla et al., 2003). Indeed the relaxation of the critical sampling constraint provides high degrees of freedom in the construction of the multiresolution scheme and it is worth taking advantage of this additional flexibility for choosing the filter shape, the bandwidths, the degree of overlapping between filters, the number of orientations, or for choosing complex-valued filters, etc. We propose here a transform optimized not only on mathematical and image processing criteria but also on biological vision, perceptual quality and natural image statistic considerations. The rationale for taking into account those additional criteria is: (1) the number of free parameters on the filter shape and multiresolution arrangement is prohibitive for a systematic study. (2) The Primary Visual Cortex supposedly evolved towards an efficient and robust image processing system adapted to natural images. Considering the important similarities between V1 simple cell characteristics and oriented multiresolution transforms, we hypothesize that biological knowledge on V1 could serve as an useful guide for the choice of the free parameters. Moreover as an interesting feedback, the implementation of biologically inspired wavelet transforms could help for the understanding of V1. (3) It is important to provide good image quality from the human perceptual point of view because digital images are almost exclusively used by humans. It would then be of great help to construct methods limiting the appearance of perceptually salient artifacts. The similarity with the biological models could serve such objective. (4) The multiresolution must be built according to the type of signal to analyse. Here as we focus on natural images, studies on their statistics can also be used as guidelines. Moreover it is worth noting that there are direct relationships between receptive fields of V1 and statistics of natural images (Field, 1987; Olshausen and Field, 1996; Doi and Lewicki, 2005).

Thus, we propose here an implementation of a multiresolution transform fulfilling the following constraints: (1) an optimal localization in space, frequency and orientation through the use of Gabor filters; (2) an augmented number of orientations; (3) a resemblance to the receptive field of V1 simple cells; (4) an exact reconstruction by a self-invertible transform (the same transform functions are used for the direct and inverse transforms) and (5) complex-valued filters sensitive to both symmetric and antisymmetric features. We pay special attention to the different aspects of the multiresolution design and we propose a series of solutions for improving the efficiency of the Gabor wavelets scheme. In comparison with previous implementations of Gabor multiresolution (Gross and Koch, 1995; Nestares et al., 1998; Ro et al., 2001; Portilla et al., 1996), the most important novelties of the present work consists in incorporating complex-valued

oriented high-pass filters, and achieving the exact reconstruction and self-invertibility properties.

The present paper fulfills two objectives: first to describe in detail in the Section 2 the proposed self-invertible log-Gabor wavelet transform which we already employed successfully in a variety of applications such as denoising (Fischer et al., 2005; Redondo et al., 2005), edge extraction (Fischer et al., 2005, in press), image compression (Fischer et al., 2006, 2005, in press) and image fusion (Redondo et al., in press), but whose implementation was incomplete or not described in detail in those publications. Second to illustrate (in Section 3) the efficiency of the method in an image denoising scenario.

2. Method

Let's consider a $n \times n$ square image \mathbf{x} of $N = n^2$ pixels. The Gabor wavelets \mathcal{W} consist in filtering the input image $\mathbf{x} \in \mathbb{R}^N$ by a set of filters $(\mathbf{G}_r)_{r \in \Omega}$. The band-pass, high-pass and low-pass filters are described respectively in Sections 2.1, 2.2 and 2.3. Direct and inverse transforms are defined in Sections 2.4 and 2.5, respectively. Section 2.6 describes the downsampling method and Section 2.7 deals with the matrix notation. Methods for obtaining exact reconstruction are described in Section 2.8.

2.1. Bandpass Log-Gabor Filters

Classical Gabor filters give rise to important difficulties when implemented in multiresolution. Filters overlap more importantly in the low frequencies than in the higher ones yielding a non-uniform coverage of the Fourier domain. Moreover Gabor filters have not zero mean, they are then affected by DC components. For those reasons log-Gabor filters are used in the present implementation instead of Gabor filters. The log-Gabor filters lack DC components and can yield a fairly uniform coverage of the frequency domain in an octave scale multiresolution scheme (Field, 1987). The log-Gabor filters are defined in the log-polar coordinates of the Fourier domain as Gaussians shifted from the origin:

$$\begin{aligned} & \mathbf{G}_{(s,t)}(\rho, \theta) \\ &= \exp\left(-\frac{1}{2} \left(\frac{\rho - \rho_s}{\sigma_\rho}\right)^2\right) \exp\left(-\frac{1}{2} \left(\frac{\theta - \theta_{(s,t)}}{\sigma_\theta}\right)^2\right) \\ & \text{with } \begin{cases} \rho_s = \log_2(n) - s & \text{if } s \text{ is odd} \\ \theta_{(s,t)} = \begin{cases} \frac{\pi}{n_r} t & \text{if } s \text{ is odd} \\ \frac{\pi}{n_r} (t + \frac{1}{2}) & \text{if } s \text{ is even} \end{cases} \\ (\sigma_\rho, \sigma_\theta) = 0.996 \left(\sqrt{\frac{2}{3}}, \frac{1}{\sqrt{2}} \frac{\pi}{n_r}\right) \end{cases} \end{aligned} \quad (1) \end{aligned}$$

where (ρ, θ) are the log-polar coordinates (in \log_2 scale, indicating the filters are organized in octave scales);

$n_s = 5$ is the number of scales of the multiresolution scheme and n_r is the number of orientations (n_r will range between 3 to 20. 8 orientations will be used as a typical value). $s \in \{1, \dots, n_s\}$ and $t \in \{1, \dots, n_r\}$ indexes the scale and the orientation of the filter, respectively; $(\rho_s, \theta_{(s,t)})$ are the coordinates of the center of the filter; and $(\sigma_\rho, \sigma_\theta)$ are the bandwidths in ρ and θ , common for all filters.

It is highly justified here to consider the Fourier domain through the log-polar coordinates. Indeed in such coordinate system, the octave distribution of the filters constitutes a regular grid, the center of the filters defined by Eq. (1) laying on a uniform hexagonal lattice. Moreover in such coordinates the filters are purely Gaussian and their bandwidths are constant across orientations and scales.

The filters defined by Eq. (1) cover only one side of the Fourier domain (see Fig 1.(a)). It is worth considering them in Fourier as the sum of a symmetric component and an antisymmetric one. Thus, those components sum their amplitude in one side of the Fourier domain and cancel themselves in the other side. This explains that in the spatial domain, the filters have both a real part (with cosine shape due to the symmetric component) and an imaginary part (with sine shape due to the antisymmetric component). A single log-Gabor filter defined in Fourier by Eq. (1) yields then both a real and an imaginary part in the spatial domain. Both parts can be seen for the 5 scale 8 orientation scheme in Fig. 1.(b) and (c), respectively. They can be compared with the filters of other multiresolution transforms (orthogonal wavelets 'Db4' (Daubechies, 1992) in Fig. 1.(e) and steerable pyramids (Portilla et al., 2003) in Fig. 1.(f)).

One objective of this study is to choose the transform parameters as close as possible of the known physiology of simple cortical cells. According to such purpose, log-Gabor filters are chosen for modeling the receptive field of the simple cortical cells as proposed in Daugman (1985); Marcelja (1980); and Field (1987). Simple cells are known to be organized in pairs in quadrature of phase (Pollen and Ronner, 1981), justifying the choice of complex-valued filters. The choice of the bandwidth in orientation is motivated by the simple cell orientation resolution which has been evaluated as around 20–40 degrees of full bandwidth at half response (Hubel, 1988; Daugman, 1985). Such orientation bandwidth would require around 6 to 13 orientations to cover the 180 degrees of the plane. For the proposed scheme using the typical value of 8 orientations we obtain filters with 31.8 degree full bandwidth at half response. The bandwidth in scale of the simple cells has been evaluated between 0.6 and 3.2 octaves (DeValois et al., 1982), and around 1.3 octaves in mean (Daugman, 1985) (at half response). The present filters have a 1.43 octave bandwidth. DeValois et al.

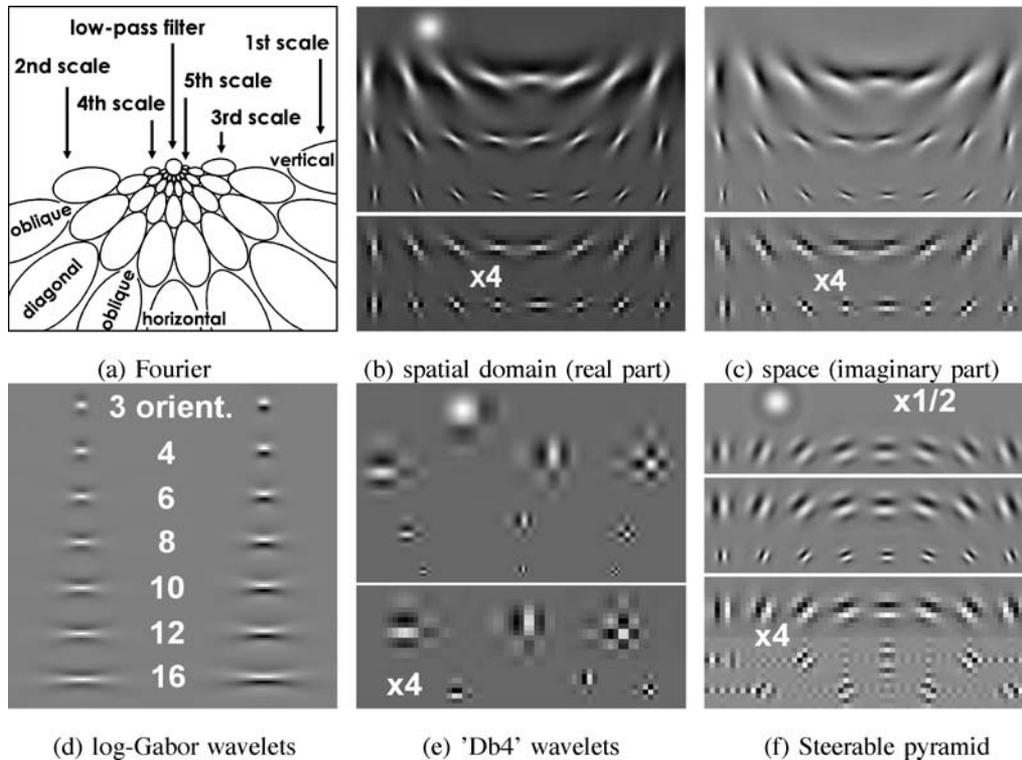


Figure 1. Multiresolution schemes. (a) Schematic contours of the log-Gabor filters in the Fourier domain with 5 scales and 8 orientations (only the contours at 78% of the filter maximum are drawn). (b) The real part of the corresponding filters is drawn in the spatial domain. The two first scales are drawn at the bottom magnified by a factor of 4 for a better visualization. The different scales are arranged in rows and the orientations in columns. The low-pass filter is drawn in the upper-left part. (c) The corresponding imaginary parts of the filters are shown in the same arrangement. Note that the low-pass filter does not have imaginary part. Insets (b) and (c) show the final filters built through all the processes described in Section 2. (d) In the proposed scheme the elongation of log-Gabor wavelets increases with the number of orientations n_r . Here the real parts (left column) and imaginary parts (right column) are drawn for the 3, 4, 6, 8, 10, 12 and 16 orientation schemes. (e) As a comparison orthogonal wavelet filters 'Db4' are shown. Horizontal, vertical and diagonal wavelets are arranged on columns (low-pass on top). (f) As a second comparison, steerable pyramid filters (Portilla et al., 2003) are shown. The arrangement over scales and orientations is the same as for the log-Gabor scheme.

(1982) reported the existence of cells covering different scales over at least 4 octaves for each retina location. At least 3 scales for each retina location have been encountered, which justifies the choice of a multiscale transform. Additionally the shape of the filters shown Fig. 1.(b)–(c) appears close to the independent components learned by sparse coding or ICA techniques on natural images (Olshausen and Field, 1996; Doi and Lewicki, 2005), confirming the adequacy of the filters to match natural image features. By its parts the modulus operation on real and imaginary parts which will be exploited for denoising in Section 3 and the resulting independency toward phase are characteristic of the complex cortical cells (Carandini et al., 2005). The thresholding operation used in the denoising method has also a biological justification (Perrinet et al., 2004). The present scheme gets then close to the V1 physiology particularly in terms of receptive field structures and multiscale and multiorientation arrangement.

2.2. High-pass Oriented Filters

An important problem appears in the first scales where from Eq. (1), \mathbf{G}_r would have significant amplitude above the Nyquist frequency ($\rho \geq \log_2(\frac{n}{2})$). Cutting off abruptly the filter response above the Nyquist frequency strongly distorts the filter shape in the spatial domain (this causes the appearance of side lobes or ringing). For this reason in many implementations high frequencies are not covered, e.g. in Ro et al. (2001). Alternatively, to not discard this part of the spectrum Nestares et al. included a non-oriented high-pass filter (Nestares et al., 1998). Nevertheless when it is used in denoising or compression, a non-oriented high-pass filter introduces cross-shaped artifacts, which are very salient and artificial-looking. We propose here several solutions to design oriented high-pass filters having smooth shape without extra side lobes. They consist first in incorporating a half-pixel shift in the spatial position of the

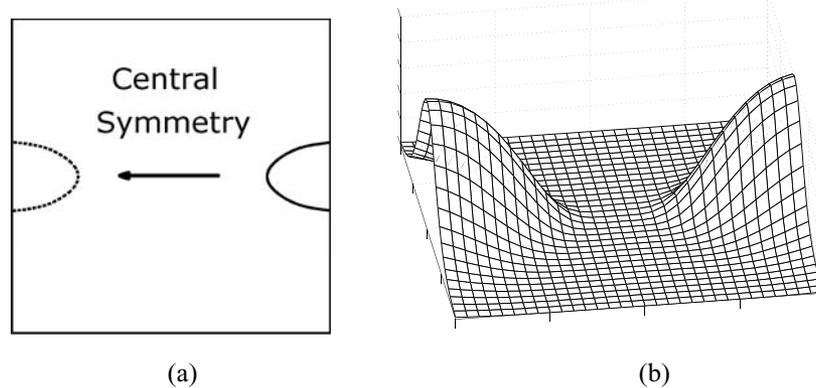


Figure 2. The real part of the first scale vertical filter is constructed in Fourier by symmetry. (a) The contours of the filters are drawn in the Fourier domain at 78% of the filter maximum. The filter is build up by summing its central symmetry. (b) Resulting filter shown in the Fourier domain (the horizontal axis represent the Fourier space which origin is located in the center of the plane and the vertical axis represents the amplitude). This filter is shown in spatial domain in Fig. 1.(b) bottom row, left column.

imaginary part of the filters. This shift allows the first scale filters to capture much more adequately the anti-symmetric features, as it will be illustrated Section 2.2.3, but consequently the real and imaginary parts of the filters have to be separately defined. Also vertical and horizontal filters are designed in a different way than the other filters, i.e. the oblique ones, see Section 2.2.1 and 2. Note that, because the definition of the first scale filters differs from the other scales, the transform is no more strictly a "wavelet" one, but it can be considered as a wavelet-like transform since the general shape of the basis functions is obtained by translation, dilatation, and rotation of a mother function if we except the modifications proposed here for improving the reconstruction performance. Moreover the proposed Gabor wavelets build directly on the original Morlet and coworkers proposition of using Gabor functions for a cycle-octave multiresolution (Morlet et al., 1982).

2.2.1. Real Part of the Horizontal and Vertical Filters.

The real part of horizontal and vertical filters is defined by central symmetry as shown in Fig. 2. Thus they are continuous across the periodicities of the Fourier domain which is important since strong discontinuities in the Fourier domain create side lobes and a worst localization in the spatial domain. They are also well localized and without extra side lobes in the space domain (see Fig. 1.(b), bottom row, 1st and 4th columns).

2.2.2. Real Part of Oblique Filters. The real part of oblique filters (i.e. filters which are not vertical nor horizontal) is also defined by central symmetry. But this is not sufficient to maintain the Fourier domain continuous (across periods) and to keep a good localization in the space domain. We propose then to fold up those filters by considering them as periodic with periodicity n . Nev-

ertheless when the part beyond the Nyquist frequency ($|u| \geq \frac{n}{2}$ or $|v| \geq \frac{n}{2}$, where (u, v) are the Cartesian coordinates of the Fourier domain) is folded up by periodicity, it covers the whole spectrum with significant amplitude. To maintain the filter selectivity to high-frequencies, it is then necessary to filter down the induced lowest frequencies by multiplying the folded part by an attenuation factor α defined here as a raised cosine function (see also Fig. 3 for an illustration of the construction of such filters):

$$\alpha = \begin{cases} \frac{1}{2} \cos\left(\pi \frac{d}{n/3}\right) + \frac{1}{2} & \text{if } d < n/3 \\ 0 & \text{if } d > n/3 \end{cases} \quad (2)$$

where d is the Euclidian distance to the closest frequency inside the Nyquist range (i.e. u and v within $[-n/2, n/2]$).

2.2.3. Imaginary Part of the Horizontal and Vertical Filters.

The imaginary part is build in Fourier similarly to the real part (Section 2.2.1) but antisymmetrically instead of symmetrically. In the spatial domain the imaginary part have to be antisymmetric, but here the axis of central antisymmetry will not be localized in a pixel but in between two pixels, this for two related reasons. First, in the spatial domain, an antisymmetric filter would have a zero as central coefficient. And, as first scale filters have a frequency close to the Nyquist frequency, i.e. an approximately two pixel period, the filter would be almost null in its central part. The filter would then be very coarsely localized, having most of its amplitude far from its center. Consequently the filter must be centered in between two pixels (see Fig. 4.(a)). Second, in the Fourier domain, an antisymmetric high-pass

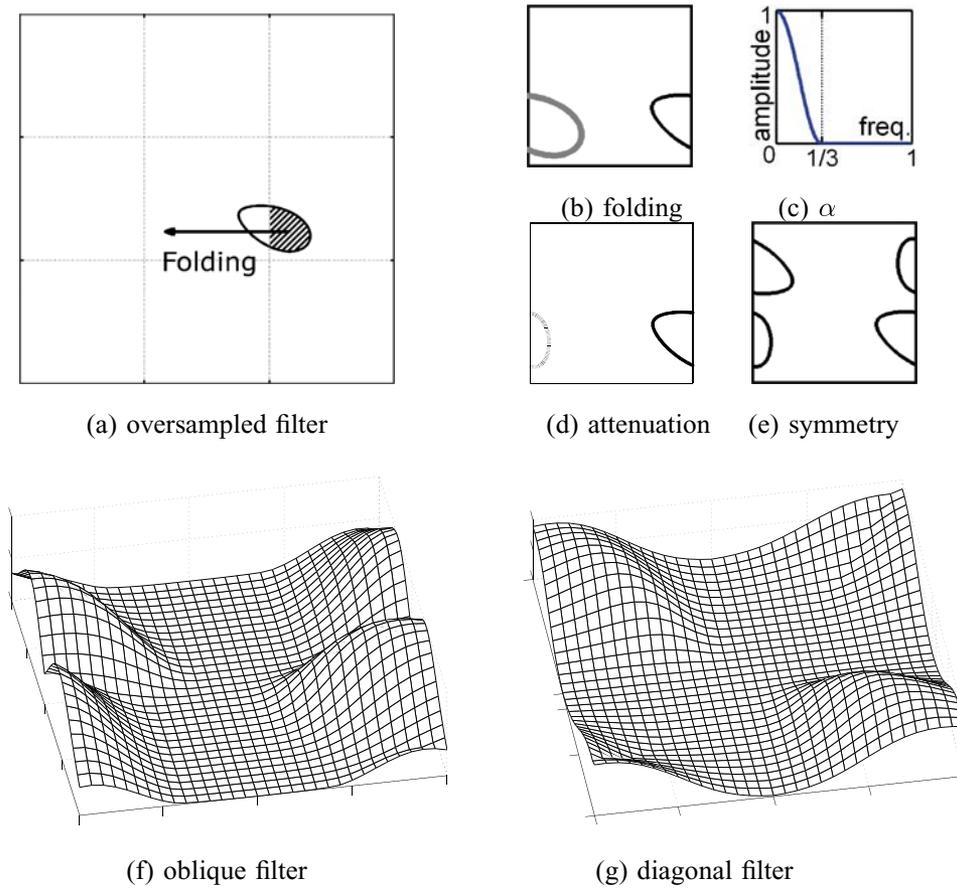


Figure 3. Construction of the real part of oblique first scale filters in the Fourier domain. (a) The frequency domain is oversampled 3 times. (b) The frequencies above the Nyquist frequency are folded up by periodicity. (c) Raised cosine function α . (d) The folded part of the filter is attenuated multiplying by the raised cosine α . (e) The central symmetry is summed up for the construction of a real-valued filter. (f) Resulting oblique filter. (g) Resulting diagonal filter. These filters are shown in the spatial domain in Fig. 1.(b) bottom row.

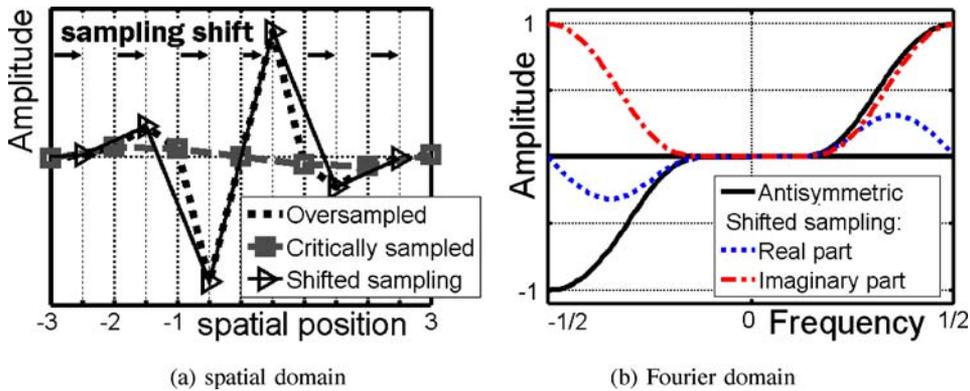


Figure 4. Imaginary part construction of horizontal and vertical first scale filters. (a) In the spatial domain it is necessary to shift the sampling grid by half a pixel. A non-shifted sampling yields a function almost null in the central coefficients, most of the amplitude being localized far from the center inducing a very poor localization of the filter. The same antisymmetric filter built on a half-pixel shifted grid is much more closely localized. (b) Shifting the sampling grid is realized by multiplying the frequencies by $e^{i\pi \frac{x}{n}}$ (inducing the filter coefficients are complex-valued also in the Fourier domain). It makes the filter continuous in its frequency coverage across periods.

filter is not continuous across periods while a half-pixel shifted version is (see Fig. 4.(b)). For those reasons the filter is multiplied by a $e^{i\pi \frac{u}{n}}$ or $e^{i\pi \frac{v}{n}}$ function in the Fourier domain (respectively for horizontal and vertical filters) which induces a half pixel displacement in the space domain. As a consequence the Fourier coefficients are now complex-valued and both the real and imaginary Fourier parts of the filters become continuous across periods (see Fig. 4.(b)).

2.2.4. Imaginary Part of Oblique Filters. Those filters are defined antisymmetrically and are folded by periodicity with the same attenuation function α described in Section 2.2.2. Moreover they are also shifted half the interpixel distance perpendicularly to the preferred direction of the filter by multiplication in Fourier by $e^{\frac{i\pi}{n}(u-t_u+v-t_v) \cdot \max(|t_u|, |t_v|)}$, where (t_u, t_v) is the normalized preferred direction vector of the filter.

2.2.5. Second Scale Filters. Second scale filters are also folded by periodicity with attenuation α (see Section 2.2.2). Here nevertheless it is not necessary to shift half-pixel in position the antisymmetric filter and the same definition is used for all orientations. An example of second scale filter is shown in Fig. 5. In the coarser scales (i.e. $s \geq 3$), the filter part beyond the Nyquist frequency has sufficiently low amplitude to be cut off.

2.3. Low-pass Filter

The multiresolution scheme is completed with a low-pass filter (DC) for recovering the luminance information. The low-pass filters could be defined simply as a Gaussian low-pass filter as in Portilla et al. (1996). Nevertheless for a better filling-in of the low-pass residual frequencies not covered by the log-Gabor filters, it is defined in a

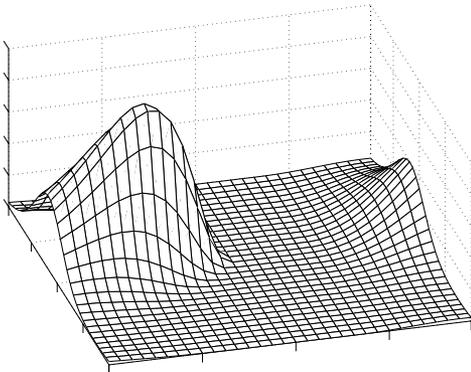


Figure 5. Second scale filters are constructed in Fourier by folding the filters by periodicity and applying the attenuation function α on the folded part. (Second scale filters are shown in the spatial domain in Fig. 1.(b),(c) next to last rows.)

different way. Two additional scales above the number of scales n_s deployed are built and summed up together (as the root sum squared). Moreover the part inside the highest additional scale is set up to one. In practice if 5 scales are deployed, the filters that would correspond to the 6th and 7th scales are summed and additionally the space inside the 7th scale ($\rho < \log_2 n - 7$) is set up to one (the resulting filter can be seen in the spatial domain in Fig. 1.(b), upper left part.)

2.4. Direct Log-Gabor Wavelet Transform

The whole set of bandpass filters, high-pass filters and the low-pass filter are indexed by $\Omega = \{1, \dots, n_s n_t + 1\}$. The filter set defined in the Fourier domain is then referred as $(\mathbf{G}_r)_{r \in \Omega}$. The convolution of the image \mathbf{x} by any filter \mathbf{G}_r is \mathbf{h}_r , which is called a *channel*. It is computed in the Fourier domain as:

$$\mathbf{h}_r = \mathcal{F}^{-1}(\mathbf{G}_r \cdot \mathcal{F}(\mathbf{x})) \quad (3)$$

where \mathcal{F} and \mathcal{F}^{-1} are respectively the direct and inverse discrete Fourier transform. The whole set of channels, $(\mathbf{h}_r)_{r \in \Omega}$, is called a *pyramid*. Finally, the log-Gabor wavelet transform of \mathbf{x} is defined by:

$$\mathcal{W}(\mathbf{x}) = (\mathbf{h}_r)_{r \in \Omega}. \quad (4)$$

2.5. Reconstruction Transform

The reconstruction transform \mathcal{W}^\dagger consists in filtering in Fourier each channel \mathbf{h}_r by the corresponding filter $\overline{\mathbf{G}}_r$ (where $\overline{\mathbf{G}}_r$ is the complex conjugate of \mathbf{G}_r . Note that in most cases the filter \mathbf{G}_r in Fourier is real-valued, so that $\overline{\mathbf{G}}_r = \mathbf{G}_r$):

$$\mathbf{z}_r = \mathcal{F}^{-1}(\overline{\mathbf{G}}_r \cdot \mathcal{F}(\mathbf{h}_r)) \quad (5)$$

Each \mathbf{z}_r is called a *reconstructed channel*. All reconstructed channels are finally summed up to obtain the reconstruction:

$$\mathcal{W}^\dagger(\mathbf{h}) = \sum_{r \in \Omega} \mathbf{z}_r = \mathcal{F}^{-1} \left(\sum_{r \in \Omega} \overline{\mathbf{G}}_r \cdot \mathcal{F}(\mathbf{h}_r) \right) \quad (6)$$

Because the image \mathbf{x} is real, its Fourier transform has hermitian symmetry. For this reason, we used band-pass Gabor filters covering only half of the Fourier plane, while the other half plane is further completed at the reconstruction using hermitian symmetry (this completion is obviated in the equations).

Because the transform is overcomplete, it can not be strictly bijective nor invertible: the transform domain has

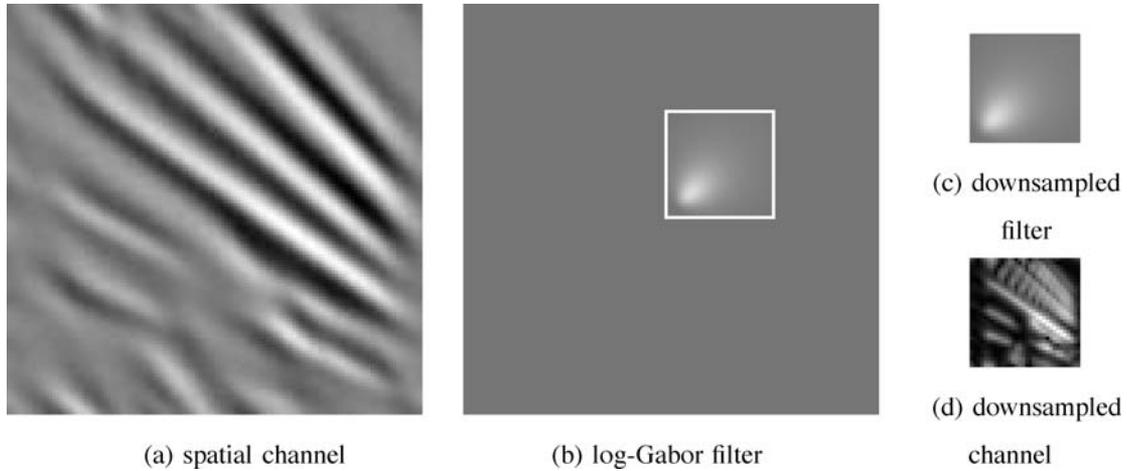


Figure 6. Sparse downsampling method. (a) 4th scale 2nd orientation channel of the image “MIT” (the original can be seen in Fig. 9). Only the moduli of the complex coefficients are shown. (b) Corresponding log-Gabor filter (scale 4, orientation 2) represented in the Fourier domain (Fourier origin is located in the center of the inset). A window of the desired downsampling size is adjusted on the filter so as to retain inside the maximum of the amplitude. Filters coefficients outside the window are zeroed out (only coefficients lower than 5% of the filter maximum are eliminated). (c) The filter can be downsampled in an invertible manner by a factor 4×4 by removing the zero coefficients. (d) The resulting downsampled channel obtained after inverse Fourier transform (as in (a) only the moduli of the coefficients are shown).

not the same dimension as the original domain. But, under the conditions that will be further exposed in Section 2.8, \mathcal{W}^\dagger can be the pseudo-inverse of \mathcal{W} so that $\mathcal{W}^\dagger(\mathcal{W}(\mathbf{x})) = \mathbf{x}$, $\forall \mathbf{x} \in \mathbb{R}^N$.

2.6. Sparse Downsampling

To limit the expansion of the transform domain it is desirable to downsample the channels as much as possible, but we desire to avoid at the same time any kind of aliasing. The *sparse downsampling* method (Wurtz, 1994) can be applied advantageously here for downsampling 3rd, 4th, 5th and low-pass filters both in u and v by factors of 2, 4, 8 and 8, respectively. The procedure is implemented in the Fourier domain as follows: a window of the defined downsampling size is centered on the filter and its position is adjusted so as to retain inside the maximum quantity of the filter amplitude (this operation is achieved by displacing iteratively the window in all the possible directions). Filter coefficients outside the window are zeroed out. In this way, only filter coefficients smaller than 5% of the filter maximum are cut off which preserves the smooth shape of the filters in the spatial domain. Using those filters, the channels can then be downsampled in an invertible manner without mixing frequencies i.e. without introducing aliasing: any frequency of the downsampled channel corresponds bijectively to one unique non-zeroed frequency in the upsampled version of the channel (the sparse downsampling procedure is illustrated in Fig. 6).

The present log-Gabor filters which are complex-valued in the spatial domain bring then an additional advantage: the filters lie only in one side of the Fourier

domain allowing the implementation of the sparse downsampling method. In comparison, pure real or pure imaginary spatial filters are symmetric or antisymmetric in Fourier and compel to downsample at Nyquist frequency, that is twice the highest frequency. The sparse downsampling compensates largely that two real values (i.e. a complex one) are obtained for each coefficient: complex-valued log-Gabor filters can still be downsampled by a factor around two times larger than real-valued filters. Consequently, the overcompleteness factor measured as the ratio M/N , where M is the dimension of the transform domain and N is the dimension of the image, is reduced to $M/N \simeq 40$ (it can be approximated as $\frac{14}{3}n_t$ where n_t is the number of orientations). It is to remark that other overcomplete schemes (Simoncelli et al., 1992; Kingsbury, 2001; Portilla et al., 2003) are able to achieve lower overcompleteness factors at the same time they maintain the translation invariance property and a relatively good directional selectivity. Nevertheless in such schemes the filters do not satisfy all the criteria defined in Section 2.1, particularly concerning the bandwidths and the filter shape. Moreover a higher overcompleteness factor is not a major drawback in many applications like denoising, the principal penalization being an increased computational cost proportional to the overcompleteness factor.

In parallel, it is straightforward to demonstrate that exact reconstruction using the same non-orthogonal filters both for decomposition and reconstruction requires the transform to be overcomplete. And to preserve the smooth shape of the filters, M must be several times larger than N . In the present case a higher downsampling would distort the filters or induce aliasing effects. The above mentioned

overcomplete factor of 40 is then necessary here. It is also relatively consistent with the redundant number of simple cortical cells compared with the number of photoreceptor neurons (Hubel, 1988) which can be evaluated around 25.

The lack of aliasing guaranties the *shiftability in space* (or *translation invariance*) (Simoncelli et al., 1992) which is of special importance for image processing tasks such as image denoising, image fusion, edge extraction, etc: the transform coefficients do not change drastically for a small displacement in space.

2.7. Matrix Notation

Because the transform is linear, it can also be viewed as a matrix operation $W \in \mathbb{R}^{N \times M}$. \mathcal{W} can then be described as the scalar product of $\mathbf{x} \in \mathbb{R}^N$ with a family of decomposition functions $(\mathbf{g}_k)_{k \in \{1, \dots, M\}}$ (the spatial log-Gabor functions), with each $\mathbf{g}_k \in \mathbb{R}^N$:

$$W\mathbf{x} = \mathbf{h} = (h_k)_{k \in \{1, \dots, M\}} = (\langle \mathbf{g}_k, \mathbf{x} \rangle)_{k \in \{1, \dots, M\}} \quad (7)$$

Because we use the same filters \mathbf{g}_k for reconstruction, the reconstruction transform is W^T (where T refers to the transposed matrix) and for all $\mathbf{h} \in \mathbb{R}^M$:

$$W^T\mathbf{h} = \sum_{k=1}^M h_k \mathbf{g}_k \quad (8)$$

Under the conditions defined in Section 2.8, W^T provides exact reconstruction. Thus the transform is said *self-invertible* (i.e. exact reconstruction is achieved using the same filters both for the decomposition and the reconstruction) (Simoncelli et al., 1992).

2.8. Methods for Reconstruction Improvement and Exact Reconstruction

The condition of exact reconstruction comes from Eq. (3) and Eq. (6) as:

$$\sum_{r \in \Omega} |\mathbf{G}_r|^2 = \mathbf{1} \quad (9)$$

Lets defined $\mathbf{S}(u, v) = \sum_{r \in \Omega} |\mathbf{G}_r|^2(u, v)$. In the sequel we propose different methods for improving the reconstruction, i.e. making \mathbf{S} closer to $\mathbf{1}$.

2.8.1. Hexagonal Fourier Lattice by Shifting Even Scales in Orientation. As already included in Eq. (1) in the definition of $\theta_{(s,t)}$, every second scale is shifted in orientation by $\frac{1}{2} \frac{\pi}{n_r}$ angle. This shift, originally proposed in Gross and Koch (1995), induces a hexagonal configuration which provides a more uniform coverage of the Fourier domain (see also Fig 1.(a).

2.8.2. Bandwidth Adjustment. The bandwidths defined by Eq. (1) have been empirically adjusted. They allow a fairly uniform coverage of the Fourier domain, the variations of \mathbf{S} measured as $2(\mathbf{S}_{\max} - \mathbf{S}_{\min})/(\mathbf{S}_{\max} + \mathbf{S}_{\min})$ being lower than 0.6% between the 2nd and 5th scale.

2.8.3. Exact Reconstruction. Each filter coefficient $\mathbf{G}_r(u, v)$ is finally normalized by $\sqrt{\mathbf{S}(u, v)}$, thus from Eq. (9) we have exact reconstruction filters. Because \mathbf{S} is already very close to the flat response before the division, the deformation introduced is small and does not produce appreciable distortions in the spatial domain (in particular side lobes are sufficiently small in amplitude to not be perceptible). We have now $\forall \mathbf{x} \in \mathbb{R}^N$, $W^T W\mathbf{x} = \mathbf{x}$. Note nevertheless that, due to the overcompleteness, in general $W W^T \mathbf{h}$ is not equal to \mathbf{h} : $\exists \mathbf{h} \in \mathbb{R}^M$, $W W^T \mathbf{h} \neq \mathbf{h}$.

3. Application to Image Denoising

The present log-Gabor wavelet scheme has already been shown efficient in many classical image processing applications such as edge extraction (Fischer et al., 2005; in press), image fusion (Redondo et al., in press), image denoising (Fischer et al., 2005; Redondo et al., 2005) and also image compression (Fischer et al., 2006, 2005, in press). Here to illustrate the efficiency of the transform we propose to compare its ability to segregate the noise from the signal in an image denoising scenario.

Most currently used denoising methods are based on anisotropic diffusion (Tschemperlé and Deriche, 2005; Sroubek and Flusser, 2003; Hamza et al., 2002) or wavelet thresholding techniques (Donoho, 1995; Coifman and Donoho, 1995; Chang et al., 2000; Portilla et al., 2003). Wavelet or multiresolution image denoising applications usually proceed in three stages: first a transformation, then a thresholding operation and finally the inverse transform for reconstructing the image. The transform aims at describing the signal in a domain where image content (principally the contrast edges) has statistically different amplitude than the noise: edges induce high amplitude coefficients while spatially incoherent noise produces a low level of amplitude spread in all the coefficients. A basic thresholding permits then to segregate most of the signal from the noise. There exist many methods for determining the optimal threshold (Donoho, 1995; Kovese, 1999; Zhong and Cherkassky, 2000; Chang et al., 2000) (see Taswell, 2000 for a review). Moreover more elaborated methods (Chang et al., 2000; Portilla et al., 2003) also use the context (i.e. they take into account the neighborhood of each coefficient instead of solely the coefficient in order to decide if it represents signal features or for noise).

Here we aim at comparing only the efficiency of the log-Gabor wavelet transform in comparison with

other multiresolutions independently of the method for threshold determination. Therefore we will test the transform in a series of classical images where Gaussian noise has been added. A gradient descent technique is applied to find the best threshold, i.e. the one giving the highest PSNR (Peak Signal-to-Noise Ratio) given the original image. The PSNR is calculated in dB as $-20\log_{10}(\sigma_e)$ where σ_e is the Root Mean Square Error (RMSE) between the original and the denoised image. Thus, the experiment finds the best result achievable if the threshold determination technique would be optimal. It measures then the ability of each transform to separate the signal from the noise by thresholding. Note that these experiments are not only significant for image denoising but also for many other image processing tasks such as image compression or edge and feature extraction for which the possibility to code the image content through a reduced set of large coefficients is of primary importance.

For each experiment all the multiresolutions are implemented with five scales. A vector of five thresholds (one for each scale) is determined by gradient descent searching for the best PSNR to the original image. No strategy has been employed for improving the response close to the borders of the image for any multiresolution. This allows to visualize the artifacts (e.g. ringing) due to the abrupt edges in those regions.

3.1. Number of Orientations

The first experiment aims at determining the optimal number of orientations to be used with log-Gabor wavelets. The question of the number of orientations is important since (bi-) orthogonal wavelets deploy just 3 orientations whereas more recent studies claimed the necessity of 6 orientations (Ro et al., 2001; Kingsbury, 2001), 8 orientations or more (Portilla et al., 2003; Fischer et al., 2005), a variable number of orientations depending of the scale (Do and Vetterli, 2005), or even up to 64 orientations (Starck et al., 2002).

Results gathered in Fig. 7 are calculated over a set of 6 classical images each of them tested with 3 different levels of noise. They show that the denoising improves importantly between 3 and 8 orientations which confirms the necessity of using more orientations than (bi-)orthogonal wavelets do. The better results are obtained around 8 to 12 orientations but with variability, since in particular cases the best results are obtained with as few as 5 or 6 orientations or as many as 16 or 20 orientations. Visual inspection of Fig. 8 allows to see first that the number of remaining noise points decreases when the number of orientations increases particularly from 6 to 8 orientations and from 8 to 12 orientations. And secondly that 12 or 16 orientations yield elongated artifacts. Such artifacts do

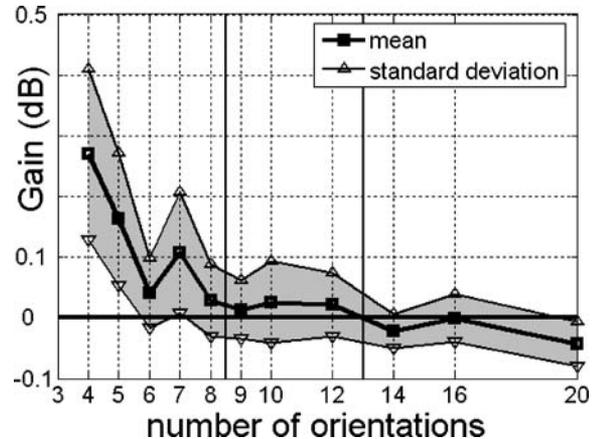


Figure 7. Denoising results as a function of the number of orientations. The denoising results are calculated over a set of 6 images (the same set as in Table 1), each of them tested with 3 different levels of noise (the PSNR is around 14, 20 and 28 dB). To reduce the influence of the large variability of conditions across experiments, the mean and standard deviation (in dB) are calculated taking as reference the previous number of orientations tested. Results show then how the denoising improves when the number of orientations is incremented from 3 to 4 orientations, from 4 to 5, from 5 to 6, etc. Results improve clearly from 3 to 8 orientations, they then achieve a maximum around 12 orientations and latter decrease slowly when more orientations are employed (see also the filter shapes as a function of the number of orientations in Fig. 1.(d)).

not appear, or at least are much less salient when using 8 orientations or less. All those results support the choice of 8 orientations as a good compromise between a low mathematical error and a high perceptual quality. This choice is also consistent with biological models of simple cells described in Section 2.1. Moreover it will allow to fairly compare with the steerable pyramids used in Portilla et al. (2003), since they also implement 8 orientations.

Note also that in the present scheme (as in other studies, see e.g. Starck et al., 2002; Do and Vetterli, 2005) the bandwidth varies depending on the number of orientations. Nevertheless it would be interesting to test both the number of orientations and the bandwidth parameters separately.

3.2. Comparison between Different Multiresolution Schemes

In Table 1 log-Gabor wavelets with 8 orientations are compared with orthogonal wavelets ‘Db4’ (Donoho, 1995), undecimated wavelets ‘Db4’ (Coifman and Donoho, 1995; Chang et al., 2000; Li and Orchard, 2000), and steerable pyramids with 8 orientations (Portilla et al., 2003; Simoncelli et al., 1992) (see the shape of the filters Fig. 1). *Soft thresholding* is applied in the first two methods. Given a threshold, it consists in the diminishing (or *shrinkage*) of all coefficients by the threshold value (and consequently it zeroes out all the coefficients which

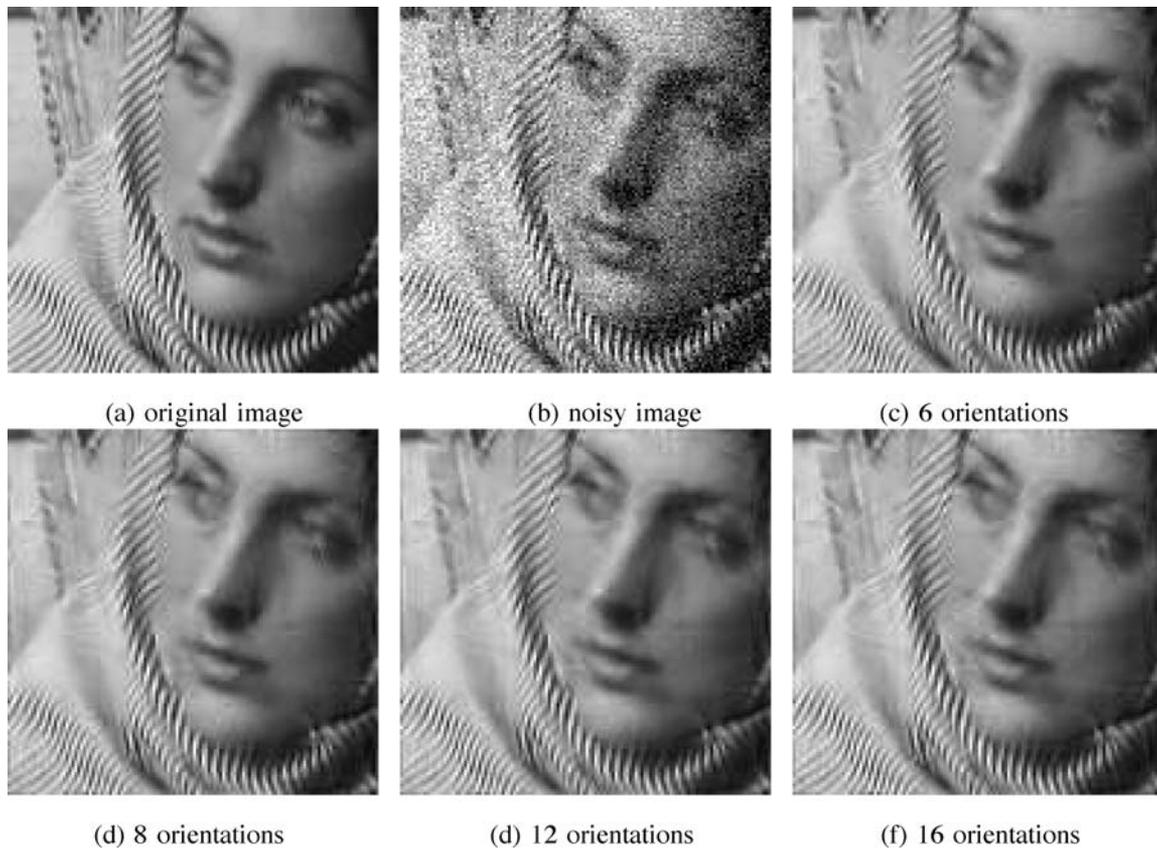


Figure 8. Denoising results through log-Gabor wavelets using different numbers of orientations. (a) 128×128 detail of the image “Barbara”. (b) The noisy image has a PSNR equal to 20.07 dB. (c) In the case of denoising using 6 orientations a significant number of high amplitude noise points are still visible. The PSNR is equal to 26.72 dB. (d) Using 8 orientations, the number of noise points is reduced. The larger size of the filters could explain in part the better segregation between edges and noise. The PSNR improves to 26.82 dB. (e) Using 12 orientations the PSNR is 27.099 dB. Almost all isolated noise points are removed but some elongated artifacts appear. (f) Although the best PSNR result is obtained here for 16 orientations (PSNR = 27.100 dB), many salient elongated artifacts appear.

amplitude lies below the threshold). Consistently with previous studies (see e.g. Donoho, 1995), we verified that for almost all the experiments soft thresholding provides better results than *hard thresholding* (i.e. zeroing out the coefficients lying below the threshold and maintaining unchanged the coefficients exceeding it) for critically sampled wavelets and also, although in a lower amount for undecimated wavelets. For orthogonal wavelets improvements are observed in all cases from the set of 6 images tested on 5 levels of noise (the same set as in Table 1). The improvement is in mean of 0.76 dB, for a standard deviation of 0.35 dB. For undecimated wavelets improvements are observed in 23 of the 30 cases and is of 0.21 ± 0.29 dB.

On the contrary hard thresholding yields better results for steerable pyramids and log-Gabor wavelets (hard thresholding is better for steerable pyramid in 25 of the 30 cases for a mean improvement of 0.24 ± 0.25 dB, and for log-Gabor wavelets in 28 of 30 cases, the improvement being in this case of 0.45 ± 0.32 dB. Full data are available in Fischer (2007)). Hard thresholding is then

used for both those last methods. For log-Gabor wavelets the threshold is applied in the modulus of the complex coefficients with exception of the first scale where it is applied separately on the real and imaginary parts of the coefficients because, as exposed in Section 2.2, they do not correspond exactly to the same position. It is not totally surprising that hard thresholding performs better here than soft thresholding since it preserves better the high amplitude coefficients which are the ones representing the image features. On the contrary the fact that soft thresholding provides better results for orthogonal and undecimated wavelets could be due to the lower degree of matching between decomposition functions and image features. This would induce that each image feature is represented by a sum of transform coefficients having very different amplitude and being consequently less adequate for the segregation by hard thresholding. The critical sampling could also enhance this effect.

It appears from Table 1 that apart from very few exceptions, undecimated wavelets always provide better results

Table 1. Denoising results given as PSNR in dB. The mean μ and standard deviation σ are given in each case in comparison with the previous method.

Original image	Noise level	Orthogonal wavelets	Undecimated wavelets	Steerable pyramids	Log-Gabor wavelets
Barbara	34.10	35.42	35.88	36.54	37.12
	28.11	30.36	30.99	32.16	32.57
	20.25	24.6	25.27	26.70	26.72
	14.74	21.08	21.64	22.69	22.92
	10.25	18.52	18.82	19.14	19.28
Teapot	34.33	35.28	35.83	36.01	36.45
	28.41	30.12	30.90	31.24	31.63
	20.71	24.44	25.28	25.62	25.79
	15.06	20.75	21.45	21.69	21.77
	10.41	17.48	17.91	17.96	18.06
Mandrill	34.12	34.29	34.35	34.30	34.35
	28.09	28.65	28.83	28.96	29.06
	20.3	22.24	22.77	23.26	23.42
	14.75	18.64	19.25	19.73	20.01
	10.28	16.22	16.67	16.98	17.16
House	34.15	35.66	36.13	36.09	36.52
	28.2	31.31	32.02	32.36	32.94
	20.38	26.86	27.59	28.28	29.10
	14.78	23.58	24.09	24.83	25.37
	10.2	19.94	20.27	20.72	20.89
Lena	33.97	34.83	35.10	35.01	35.45
	28.15	29.79	30.23	30.54	31.11
	20.34	24.37	24.99	25.52	26.01
	14.89	21.27	21.80	22.16	22.59
	10.26	18.32	18.67	18.78	18.94
Mit	34.82	35.73	36.12	36.06	36.94
	28.79	30.51	31.10	31.22	32.29
	20.98	24.10	24.85	25.29	25.95
	15.56	20.00	20.57	21.25	21.77
	10.64	16.21	16.56	16.86	17.10
Mean	34	+0.95 \pm 0.42	+0.37 \pm 0.16	+0.10 \pm 0.26	+0.47 \pm 0.25
improvement	28	+1.83 \pm 0.76	+0.55 \pm 0.20	+0.40 \pm 0.35	+0.52 \pm 0.29
($\mu \pm \sigma$)	20	+3.94 \pm 1.38	+0.69 \pm 0.096	+0.65 \pm 0.36	+0.39 \pm 0.29
	15	+5.92 \pm 1.58	+0.58 \pm 0.062	+0.59 \pm 0.27	+0.35 \pm 0.16
	10	+7.44 \pm 1.43	+0.37 \pm 0.053	+0.26 \pm 0.14	+0.16 \pm 0.044

than critically sampled wavelets (+0.51 dB of improvement in mean, the standard deviation being $\sigma = 0.18$ dB). Steerable pyramids provide better results than undecimated wavelets (+0.40 \pm 0.36 dB of improvement) and log-Gabor wavelets outperform all the three former methods (+0.38 \pm 0.26 dB of improvement compared with steerable pyramids. Improvements in terms of PSNR are observed in each of the 30 cases tested which implies the improvement is statistically highly significant. It is to note moreover that the diversity of images and of noise levels cover a wide range of denoising situations). Vi-

sual inspection on the Fig. 9 and 10 show that denoised images contain artifacts which shape is characteristic of the multiresolution employed. Indeed the shape of the artifacts corresponds to the shape of the decomposition functions shown Fig. 1. Residual features due to noise also remain, their quantity could be reduced by a more severe thresholding. Here the thresholds are adjusted for the best PSNR. The thresholding should be a bit higher for avoiding most of the residual noise at the cost of losing an additional part of the signal. This would be perceptually more pleasant even if the PSNR would be

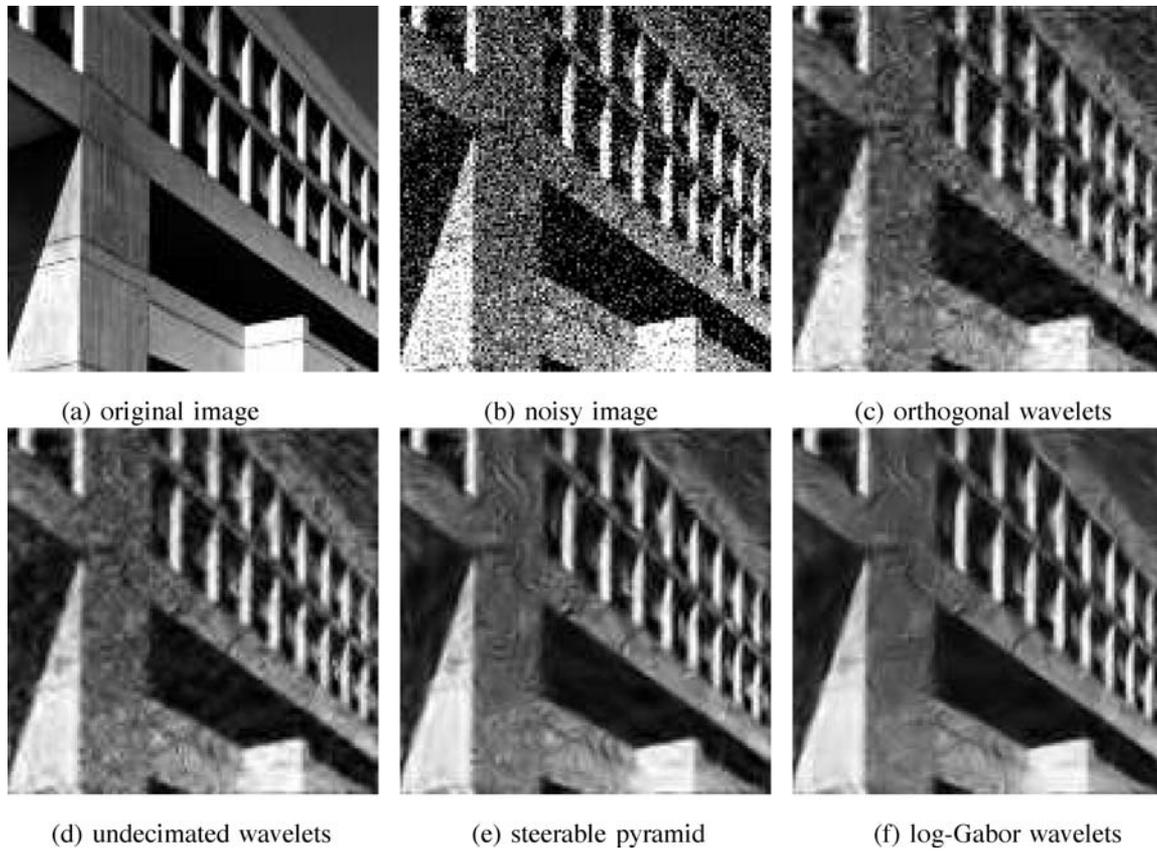


Figure 9. Denoising results using different multiresolution schemes. (a) 128×128 pixel detail of the image “MIT”. (b) The noisy version has a PSNR of 15.56 dB. (c) Denoising by orthogonal wavelets provides a 20.00 dB denoised image. (d) Undecimated wavelets yield a 20.57 dB denoised image. (e) Steerable pyramid yields a 21.25 dB denoised image. (f) Log-Gabor wavelets achieve a 21.77 dB denoised image (see also the entire result set in Table 1).

worst. It is remarkable that the log-Gabor method generally induces less amount of artifacts (particularly when compared with orthogonal and undecimated wavelets but also in comparison with steerable pyramids) which would indicate an augmented statistical separation in the transform domain between edge and noise features. In addition the artifacts can appear more natural looking so that they could result less annoying and be more easily perceived as image features (specially for images without human features, see Fig. 9).

It has already been established that the translation invariance property improves the results of overcomplete representations in comparison to critically sampled ones (Simoncelli et al., 1992; Coifman and Donoho, 1995; Chang et al., 2000; Li and Orchard, 2000), also because of the aliasing present in critically sampled wavelets. The better resolution in orientation can explain the better results of steerable pyramids and log-Gabor wavelets on undecimated wavelets. Note moreover that the finer bandwidth in orientation yields larger filters in the space domain, consequently it involves more pixels in the calculation of the coefficients: the noise is averaged

on more pixels what leads to a better noise robustness (Fig. 1.(d) shows how the filter size increases with the number of orientations).

The better performances of the log-Gabor representation in comparison to steerable pyramids could be explained by four factors: (1) the Gabor (Gaussian) filter shape offers an optimal joint localization in frequency and space, which is an improvement in comparison with the raised cosine shape of steerable pyramids. (2) Log-Gabor functions are complex-valued with odd and even parts which permit a better capture of both edges and ridges (Kovesi, 1999; Fischer et al., 2005). (3) The oriented high-pass log-Gabor filters are smooth and without extra side-lobes in space (while the high-pass steerable filters (Portilla et al., 2003) shown in Fig. 1.(f), bottom row seem to be constructed without the precautions described in Section 2.2). (4) The proposed log-Gabor filters have elongated shape whereas steerable filters are more isotropic in size, i.e. log-Gabors have larger bandwidth in frequency (1.43 octave against 1 octave for steerable pyramids) and narrower bandwidth in orientation (37 degrees against 50 degrees for

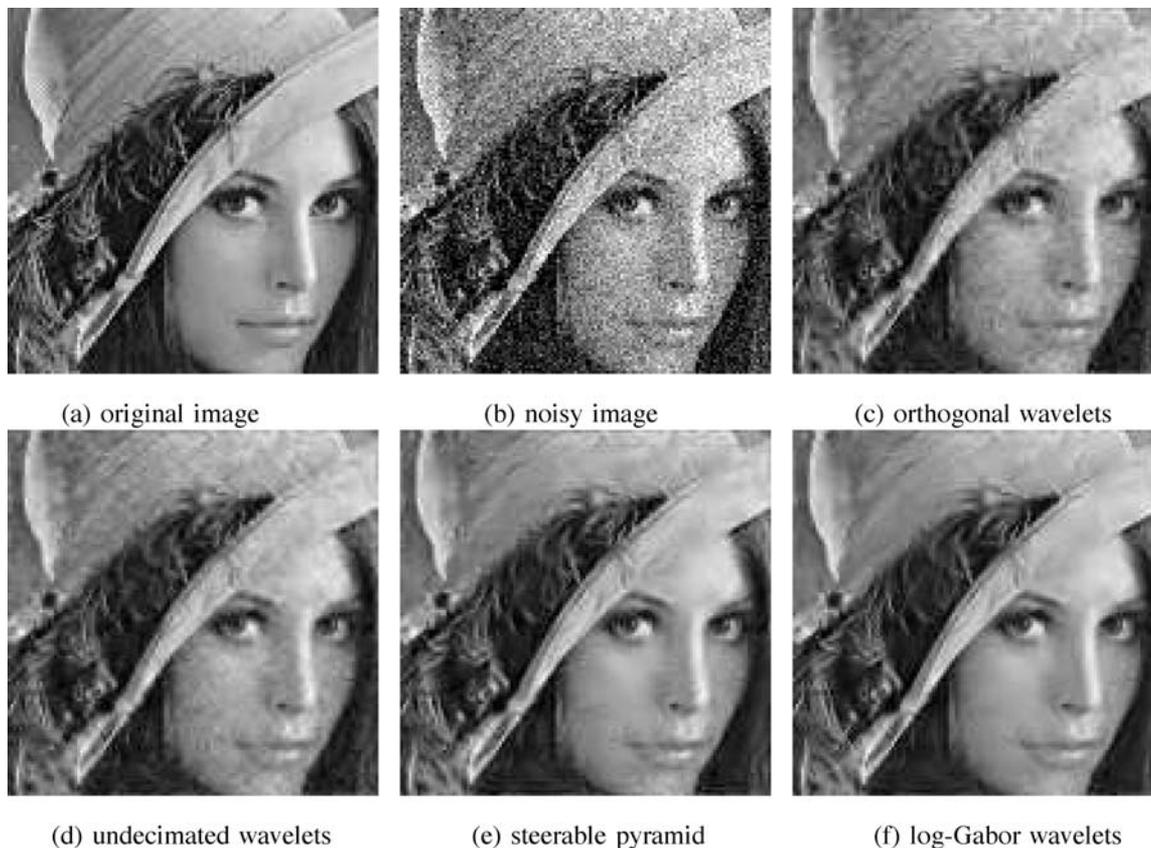


Figure 10. Denoising results for different multiresolution schemes. (a) 128×128 detail of the image “Lena”. (b) The noisy version has a PSNR equal to 20.34dB. (c) Denoising by orthogonal wavelets provides a 24.37 dB denoised image. (d) Undecimated wavelets yield a 24.99 dB denoised image. (e) Steerable pyramid yields a 25.52 dB denoised image. (f) Log-Gabor achieves a 26.01 dB denoised image. (see the entire result set in Table 1).

steerable). The elongated shape seems more appropriate for two reasons: they are adapted from biological data and they are also closer to the independent component of the natural images extracted by sparse coding or ICA techniques (Olshausen and Field, 1996; Doi and Lewicki, 2005). Thus the log-Gabor functions should match better with edges of natural images (yielding a stronger statistical differential response between edge and noise features) and as an additional advantage they can appear more “natural looking” to human observers.

4. Conclusions

We proposed an overcomplete multiresolution scheme mimicking the receptive field properties of simple cortical cells and optimized for achieving exact reconstruction through carefully designed filters. The proposed log-Gabor wavelet transform is optimized by taking into account mathematical, biological vision and image statistic considerations. From the mathematical point of view, Gabor functions are optimally joint-localized in frequency and space, which makes them optimal functions

to characterize signals. The log-Gabor wavelet transform is moreover self-invertible which has been shown important for preventing the appearance of artifacts. Concerning biological vision criteria, the log-Gabor filters mimic closely the receptive field of V1 simple cortical cells. The proximity to biological vision can help simultaneously for choosing adequate transform parameters, for limiting the saliency of the artifacts and for further developing the biological models. And third, in relation with statistics of natural images, the transform filters have similar shape as the independent components learned from natural images, which support the proposed log-Gabor wavelets as an adequate scheme for matching natural image features.

The transform showed excellent results for segregating the signal from the noise by hard thresholding. The comparison with other methods (orthogonal wavelets, undecimated wavelets and steerable pyramids) showed an overall better performance of the proposed log-Gabor wavelet technique both in terms of the mathematical error and perceptual quality. The transform confirms then to be an adequate tool for representing features of natural images and to segregate them from incoherent noise.

Among further developments, a whole image denoising application could be set up by implementing a dedicated threshold determination method which would take into account the relationships between neighboring coefficients as proposed e.g. in Chang et al. (2000); Portilla et al. (2003). Many other applications can also be derived since the ability of coding image features through a reduced set of large transform coefficients makes the transformation suitable for many applications such as image fusion or feature extraction. For image compression but also for general applications the redundancy induced by the overcompleteness could be recovered by deploying sparse approximation algorithms (Perrinet, 2004; Fischer et al., in press).

Acknowledgments

This work has been supported in part by the grants TEC 2004-00834, PI040765, TEC2005-24046-E, TEC 2005-24739-E and the bilateral project 2004CZ0009 of CSIC and Academy of Sciences of the Czech Rep. SF, RR and LP are supported by grants from MEC-FPU, CSIC-I3P and EC IP project FP6-015879, "FACETS", respectively. FS is supported by grants of the Spanish States Secretary of Education and Universities fellowship, and grants No. 102/04/0155, No. 202/05/0242 of the Grant Agency of the Czech Republic.

References

- Carandini, M., Demb, J.B., Mante, V., Tolhurst, D.J., Dan, Y., Olshausen, B.A., Gallant, J.L., and Rust, N.C. 2005. Do we know what the early visual system does? *J Neurosci*, 25(46):10577–10597.
- Chang, S.G., Yu, B., and Vetterli, M. 2000. Adaptive wavelet thresholding for image denoising and compression. *IEEE Trans. on Image Proc.*, 9(9):1532–1546.
- Chang, S.G., Yu, B., and Vetterli, M. 2000. Spatially adaptive wavelet thresholding with context modeling for image denoising. *IEEE Trans. on Image Processing*, 9(9):1522–1531.
- Christiansen, O. 2002. *Time-Frequency Analysis and its Applications in Denoising*. PhD thesis, Department of Informatics, University of Bergen, Norway.
- Clausi, D.A. and Jernigan, M.E. 2000. Designing Gabor filters for optimal texture separability. *Pattern Recognition*, 33:1835–1849.
- Coifman, R.R. and Donoho, D. 1995. Translation-invariant de-noising. In *Wavelets and Statistics, Lecture Notes in Statistics* 103, Springer Verlag: NY. A. Antoniadis and G. Oppenheim (eds), pp. 125–150.
- Cristóbal, G. and Navarro, R. 1994. Space and frequency variant image enhancement based on a Gabor representation. *Patt. Rec. Letters*, 15(3):273–277.
- Daubechies, I. 1992. *Ten Lectures on Wavelets*. SIAM, Philadelphia: PA.
- Daugman, J. 1985. 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.
- Daugman, J. 1988. Complete discrete 2-D Gabor transforms by neural networks for image analysis and compression. *IEEE Trans. Acoust. Speech Signal Proc.*, 36(7):1169–1179.
- DeValois, R.L., Albrecht, D.G., and Thorell, L.G. 1982. Spatial frequency selectivity of cells in macaque visual cortex. *Vision Res.*, 22:545–559.
- Do, M.N. and Vetterli, M. 2005. The contourlet transform: An efficient directional multiresolution image representation. *IEEE Trans. on Image Proc.*, 14(12):2091–2106.
- Doi, E. and Lewicki, M.S. 2005. Relations between the statistical regularities of natural images and the response properties of the early visual system. *Japanese Cognitive Science Society, SIG P&P*, pp. 1–8.
- Donoho, D. 1995. De-noising by soft-thresholding. *IEEE Trans. Inf. Theory*, 41(3):613–627.
- Field, D.J. 1987. Relation between the statistics of natural images and the response properties of cortical cells. *J. Opt. Soc. Am. A*, 4(12):2379–2394.
- Fischer, S. 2007. New contributions in overcomplete image representations inspired from the functional architecture of the primary visual cortex. PhD thesis, Technical University Madrid, High Technical School of Telecommunication Engineering, Dept. of Electronic Engineering, Spain.
- Fischer, S., Cristóbal, G., and Redondo, R. 2006. Sparse overcomplete Gabor wavelet representation based on local competitions. *IEEE Trans. on Image Proc.*, 15(2):265–272.
- Fischer, S., Redondo, R., and Cristóbal, G. 2005. Sparse edge coding using overcomplete Gabor wavelets. In *IEEE Int. Conf. on Image Proc.*, vol. 1, pp. 85–88.
- Fischer, S., Redondo, R., Perrinet, L., and Cristóbal, G. 2005. Sparse Gabor wavelets by local operations. In *Proc. SPIE, Bioengineered and Bioinspired Systems II*, G. Linan-Cembrano; R.A. Carmona (eds), Vol. 5839, pp. 75–86.
- Fischer, S., Redondo, R., Perrinet, L., and Cristóbal, G. Sparse approximation of images inspired from the functional architecture of the primary visual areas. *EURASIP JASP, special issue on Image Perception* in press.
- Gabor, D. 1946. Theory of communication. *J. Inst. Electr. Eng.*, 93:429–457.
- Grigorescu, C., Petkov, N., and Westenberg, M.A. 2003. Contour detection based on nonclassical receptive field inhibition. *IEEE Trans. on Image Proc.*, 12(7):729–739.
- Gross, M.H. and Koch, R. 1995. Visualization of multidimensional shape and texture features in laser range data using complex-valued Gabor wavelets. *IEEE Trans. Visual. and Comput. Graphics*, 1(1):44–59.
- Hamza, A.B., Krim, H., and Unal, G. 2002. Unifying probabilistic and variational estimation. *IEEE Signal Proc.*, 19(5):37–47.
- Heitger, F., Von der Heydt, R., Peterhans, E., Rosenthaler, L., and Kubler, O. 1998. Simulation of neural contour mechanisms: Representing anomalous contours. *Image and Vision Computing*, 16:409–423.
- Hubel, D. 1988. *Eye, Brain, and Vision*. WH Freeman: Sc. Am. Lib. Series, New York.
- Kingsbury, N.G. 2001. Complex wavelets for shift invariant analysis and filtering of signals. *Jour. of Applied and Comput. Harmonic Analysis*, 10(3):234–253.
- Kovesi, P. 1999. Phase preserving denoising of images. In *Australian Patt. Recog. Soc. Conf. DICTA*. Perth WA., pp. 212–217.
- Kovesi, P. 2003. Phase congruency detects corners and edges. In *Australian Patt. Recog. Soc. Conf. DICTA*. Sydney WA., pp. 309–318.
- Krüger, V. 2001. *Gabor Wavelet Networks for Object Representation*. PhD thesis, Christian-Albrechts-University Kiel, Technical Faculty.
- Lee, T.S. 1996. Image representation using 2D Gabor wavelets. *IEEE Trans. Pattern Anal. Mach. Intell.*, 18(10):959–971.
- Li, X. and Orchard, M.T. 2000. Spatially adaptive image denoising under overcomplete expansion. In *Int. Conf. on Image Proc.*, Vol. 3, pp. 300–303.

- Marcelja, S. 1980. Mathematical description of the responses of simple cortical cells. *J. Opt. Soc. Am. A*, 70(11):1297–1300.
- Mingolla, E., Ross, W., and Grossberg, S. 1999. A neural network for enhancing boundaries and surfaces in synthetic aperture radar images. *Neural Networks*, 12(3):499–511.
- Morlet, J., Arens, G., Fourgeau, E., and Girard, D. 1982. Wave propagation and sampling theory. *Geophysics*, 47:203–236.
- Nestares, O., Navarro, R., Portilla, J., and Taberner, A. 1998. Efficient spatial-domain implementation of a multiscale image representation based on Gabor functions. *Jour. of Electr. Imag.*, 7(1):166–173.
- Olshausen, B. and Field, D.J. 1996. Wavelet-like receptive fields emerge from a network that learns sparse codes for natural images. *Nature*, 381:607–609.
- Perrinet, L. 2004. Feature detection using spikes: The greedy approach. *J. Physiology (Paris)*, 98(4–6):530–539.
- Perrinet, L., Samuelides, M., and Thorpe, S. 2004. Coding static natural images using spiking event times: Do neurons cooperate? *IEEE Trans. on Neural Networks*, 15(5):1164–1175.
- Pollen, D.A. and Ronner, S.F. 1981. Phase relationships between adjacent simple cells in the visual cortex. *Science*, 212:1409–1411.
- Portilla, J., Navarro, R., Nestares, O., and Taberner, A. 1996. Texture synthesis-by-analysis based on a multiscale early-vision model. *Opt. Eng.*, 35(8):1–15.
- Portilla, J. and Simoncelli, E.P. 2000. A parametric texture model based on joint statistics of complex wavelet coefficients. *Int. Jour. Comp. Vis.*, 40(1):49–70.
- Portilla, J., Strela, V., Wainwright, M., and Simoncelli, E. 2003. Image denoising using scale mixtures of Gaussians in the wavelet domain. *IEEE Trans. Image Proc.*, 12(11):1338–1351.
- Pöttsch, M., Krüger, N., and Malsburg, C. 1996. Improving object recognition by transforming Gabor filter responses. *Network: Computation in Neural Systems*, 7(2):341–347.
- Redondo, R., Fischer, S., Perrinet, L., and Cristóbal, G. 2005. Modeling of simple cells through a sparse overcomplete Gabor wavelet representation based on local inhibition and facilitation. In *European Conf. On Visual Perception*, A Coruña, Spain.
- Redondo, R., Sroubek, F., Fischer, S., and Cristóbal, G. in press. A multiresolution-based fusion scheme through log-Gabor wavelets and a multisize windows technique. *Information Fusion*.
- Ro, Y.M., Kim, M., Kang, H.K., Manjunath, B.S., and Kim, J. 2001. MPEG-7 homogeneous texture descriptor. *ETRI Journal*, 23(2):41–51.
- Simoncelli, E.P., Freeman, W.T., and Heeger, D.J. 1992. Shiftable multiscale transforms. *IEEE Trans. Inf. Theory*, 38(2):587–607.
- Sroubek, F. and Flusser, J. 2003. Multichannel blind iterative image restoration. *IEEE Trans. Image Proc.*, 12(9):1094–1106.
- Starck, J.L., Candès, E.J., and Donoho, D.L. 2002. The curvelet transform for image denoising. *IEEE Trans. on Image Proc.*, 11(6):670–684.
- Taswell, C. 2000. The what, how and why of wavelet shrinkage denoising. *Computing in Science and Engineering*, pp. 12–19.
- Tschumperlé, D. and Deriche, R. 2005. Vector-valued image regularization with PDE's: A common framework for different applications. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 27(4):506–517.
- Wurtz, R. 1994. *Multilayer Dynamic Link Networks for Establishing Image Point Correspondences and Visual Object Recognition*. PhD thesis, Bochum University.
- Zhong, S. and Cherkassky, V. 2000. Image denoising using wavelet thresholding and model selection. In *Int. Conf. on Image Proc.*, Vol. 3, pp. 262–265.