EFFICIENT SPATIAL-DOMAIN IMPLEMENTATION OF A MULTISCALE IMAGE REPRESENTATION BASED ON GABOR FUNCTIONS

Oscar Nestares, Rafael Navarro, Javier Portilla and Antonio Tabernero*
Instituto de Optica "Daza de Valdés" (CSIC). Serrano 121. 28006 Madrid. Spain.
*Facultad de Informática. Universidad Politécnica de Madrid. Boadilla del Monte. Madrid. Spain

ABSTRACT

Gabor schemes of multiscale image representation are useful in many computer vision applications. However, the classic Gabor expansion is computationally expensive due to the lack of orthogonality of Gabor functions. Some alternative schemes, based on the application of a bank of Gabor filters, have important advantages such as computational efficiency and robustness, at the cost of redundancy and lack of completeness. In a previous work we proposed a quasicomplete Gabor transform, suitable for fast implementations in either space or frequency domains. Reconstruction was achieved by simply adding together the even Gabor channels. In this work, we develop an optimized spatial-domain implementation, using one-dimensional, 11-tap filter masks, that is faster and more flexible than Fourier implementations. The reconstruction method is improved by applying fixed and independent weights to the Gabor channels before adding them. Finally, we analyze and implement, in the spatial domain, two ways to incorporate a high-pass residual, which permits a visually complete representation of the image.

Keywords. Gabor functions, multiscale representation, spatial-domain implementation.

1. INTRODUCTION

Image representations based on Gabor functions, or other similar wavelets, have shown to be highly useful in many applications (for a review, see [1]) such as image coding and compression [2], enhancement and restoration [3], or analysis (texture analysis [4, 5] and synthesis [6], motion [7], etc.) Moreover, Gabor functions (frequently in the frame of multiscale filtering schemes) are often used in current models of image representation in the visual cortex [8-11] because they are a good approximation to the receptive fields of simple cortical cells [12, 13]. However, Gabor functions are not orthogonal and, as a consequence, the classic Gabor expansion is computationally expensive, having unusual dual basis functions [14]. The reconstruction requires the use, among others, of iterative algorithms [15], neural networks [2], or the inversion of large matrices [16]. These problems can be partially overcome by using a redundant, multiscale filtering implementation, as the one we now propose.

Among the many wavelet, multiresolution pyramids and related schemes using different basis functions (such as Gaussian derivatives [17], steerable filters [18], etc.), those based on Gabor functions (GF) have several advantages: (1) maximizing joint localization in both spatial and frequency domains [13]; (2) flexibility; GF can be freely tuned to a continuum of spatial positions, frequencies and orientations, using arbitrary bandwidths; (3) GF are the only biologically plausible filters with orientation selectivity that can be exactly expressed as a sum of only two separable filters [7]. This unique property has made possible the efficient implementation proposed in this paper; (4) GF perform well in a large variety of applications. For all the above reasons, Gabor functions are especially suitable for performing early processing tasks in multipurpose environments of image analysis and machine vision.

In a previous work we proposed a vision-inspired scheme using Gabor functions [8]. This scheme was characterized by an overcomplete representation that permitted fast direct and inverse transforms and by a log-polar sampling of the frequency domain. This log-polar sampling provided both frequency and orientation selectivity, but did not adequately cover the corners of the frequency domain, i.e., the very high frequencies along the diagonals. This area was disregarded due to three reasons: first, the power spectra of natural images tend to show an exponential decay with frequency [19], so that there is very little energy at these frequencies; secondly, in most cases the image acquisition is made through a conventional optical system, which filters out the very high frequencies. Finally, many image analysis tasks do not require considering the high frequencies region (such as those related with texture [4, 6], motion [7], etc.) We obtained a satisfactory reconstruction (with small amplitude distortion) by simply adding all the even Gabor channels, that provided a good coverage of the Fourier domain except for the very high frequencies.

In this paper, we study different solutions to optimize the spatial implementation of the Gabor scheme. We focus on the spatial domain implementation because it has two important advantages with respect to the Fourier one: (1) it is more plausible for modeling vision (the visual system does not seem to compute Fourier

transforms); and (2) it permits local processing which is restricted to areas of interest and non rectangular shapes. We propose a significantly improved spatial domain implementation with respect to the original work [8], in which: (1) we incorporate a High-Pass Residual (HPR) covering the high frequencies, in a way similar to that proposed by Watson in his cortex transform [20]. We study different strategies for implementing the HPR in the spatial-domain (contrarily to the Fourier implementation of the cortex transform) that do not significantly increase the computational cost; (2) the quality of the reconstruction is highly improved by assigning different fixed gains to the Gabor channels before adding them together; and (3) we use separable one-dimensional filter masks with small size (11-tap), resulting in a spatial domain implementation faster than the one in the frequency domain via FFT, while maintaining a high fidelity in the filter design. One of the main consequences of these optimizations is the improvement in the objective quality of the reconstructed images, that are visually indistinguishable from the original ones.

This paper is organized as follows. In Section 2 we review the previously proposed Gaussian wavelet transform, and introduce the theoretical formulation of an exact HPR. Section 3 describes the implementation details, including the design of the filters and the different formulations for the HPR. Section 4 deals with the image reconstruction methods, evaluating the objective quality achieved by the different options. In Section 5 we compare the computational efficiency of our spatial domain implementation to the Fourier domain one. Finally, a summary is presented in Section 6.

2.- GAUSSIAN WAVELET TRANSFORM

The Gaussian wavelets, or Gabor functions, are complex exponential functions modulated by Gaussian functions. In the particular case of the Gaussian envelope having circular symmetry, and assigning a zero phase to the complex exponential, the expression for a Gabor function tuned to the frequency f_0 and orientation θ_0, and centered at the origin ($x_0 = 0, y_0 = 0$) is:

$$g_{0,0,f_0,\theta_0}(x,y) = \exp(-\pi a^2(x^2 + y^2)) \cdot \exp(i2\pi f_0(x \cdot \cos\theta_0 + y \cdot \sin\theta_0)), \qquad (1)$$

where a determines the spatial frequency bandwidth. In our case, we have set $a = \left(3\sqrt{\ln 2 / \pi}\right)^{-1} f_0$, which yields a constant (in a logarithmic scale) 1 octave bandwidth in the frequency domain. The four subscripts account for localization in both the spatial and the frequency domains, (0, 0) and (f_0, θ_0), respectively. From the basic Gabor function (Eq. 1) the complete set of functions used for sampling the joint space-frequency domain is obtained by **rotations**, with a 45^0 step, to get four orientation channels in the frequency domain; by **stretching**, by a factor of two, to halve the frequency of the sinusoid (hence sampling the frequency domain in octaves); and by **translations** in x and y (with an interval $(1/4)f_0$) to sample the spatial domain. We have limited the number of spatial frequencies to four, starting with $f_4 = f_N / 2$ (f_N being the Nyquist frequency), and following with $f_3 = f_4 / 2$; and so on. Although more channels could be easily included, we have found that 4x4 channels provide good results for most applications (moreover, there is

psychophysical evidence that the visual system uses a similar number of channels [11]).

The coefficients of the Gabor transform are simply obtained through convolution of the image with each one of the filters. Figure 1a shows the even real part of the 4x4 Gabor filters of the scheme. Only the real part is necessary to represent and reconstruct the image [8], although for many applications both real and imaginary parts are useful.

############# **Insert Figure 1 about here** #############

Figure 1b represents the coverage of the Fourier domain obtained with our set of Gabor functions. If this coverage were perfectly flat, the reconstruction would consist of simply adding the channels. However, there are regions that do not have a perfect flat response, as well as regions that are not covered at all (the very low frequencies, and the very high frequencies). We have included both low-pass (LPR) and high-pass (HPR) residual channels [20] to obtain a perfect covering of the Fourier domain. For homogeneity, we have used a Gabor function for the LPR, with frequency $f_0 = 0$:

$$g_{0,0,0,0}(x,y) = \exp(-\pi a^2 (x^2 + y^2)),\tag{2}$$

where the parameter a has been set to provide a good covering of the low frequencies region (equal to that of the Gabor channel with frequency f_2). Finally, we can obtain the expression of the HPR that makes the sum of itself with all the Gabor channels (including the LPR) to give a flat response. Its impulse response in the spatial domain is:

$$h(x,y) = \delta(x,y) - g_{0,0,0,0}(x,y) - \sum_{f,\theta} \mathrm{Re}\big[g_{0,0,f,\theta}(x,y)\big].\tag{3}$$

Figure 2 represents the impulse response of this HPR. This filter would permit an exact reconstruction of the original image because, apart from containing the very high frequencies not included in the Gabor channels, it also compensates for the small amplitude distortions at lower frequencies. The main drawbacks of this filter are that it does not belong to the family of filters of the scheme (Gabor functions in our case) and that its implementation in the spatial domain is computationally expensive. In the next Section we propose and discuss different approaches to overcome the last problem.

############# **Insert Figure 2 about here** #############

3.- EFFICIENT IMPLEMENTATION OF THE GWT IN THE SPATIAL DOMAIN

Multiresolution pyramidal implementation

The distribution by octaves of the filters' tuning frequencies allows one to use an efficient multiresolution pyramidal [21] implementation of the Gabor wavelet transform. The same set of filters is applied to reduced (through low-pass filtering and down sampling in a factor of two) versions of the image. The computational

4

advantages of the multiresolution implementation of multiscale schemes are well known.

 We have applied to the images a separable cubic B-spline filter [22] before subsampling to avoid aliasing. This filter has very good properties, while having only five elements. The LPR is obtained by low-pass filtering again the 4th layer low-pass version of the original image obtained in the multiresolution pyramid. This LPR has a spectral response very similar (signal-to-noise ratio -SNR- of 25 dB) to the desired Gaussian one (Eq. 2).

Design of Gabor Filters

In designing the filters we have considered two criteria to maximize: (1) fidelity of the frequency response and (2) computational efficiency. The use of separable filter masks is a very efficient strategy to reduce computational cost, still keeping high fidelity. A Gabor function having its principal axes (in the frequency domain) parallel to the frequency axes can be exactly expressed as the sum of two separable Gabor functions [7]. The following equations show this property for both the real and imaginary parts:

$$
\begin{aligned}
R_{0,0,f_0,\theta_0,p=0}(x,y) &= \left[g_x \cdot \cos(2\pi f_0 x \cos\theta_0)\right] \cdot \left[g_y \cdot \cos(2\pi f_0 y \sin\theta_0)\right] \\
&\quad -\left[g_x \cdot \sin(2\pi f_0 x \cos\theta_0)\right] \cdot \left[g_y \cdot \sin(2\pi f_0 y \sin\theta_0)\right] \\
R_{0,0,f_0,\theta_0,p=1}(x,y) &= \left[g_x \cdot \sin(2\pi f_0 x \cos\theta_0)\right] \cdot \left[g_y \cdot \cos(2\pi f_0 y \sin\theta_0)\right] \\
&\quad +\left[g_x \cdot \cos(2\pi f_0 x \cos\theta_0)\right] \cdot \left[g_y \cdot \sin(2\pi f_0 y \sin\theta_0)\right]
\end{aligned}
\tag{4}
$$

Here g_x and g_y represent 1D Gaussians ($g_x = \exp(-\pi a x^2)$), and the fifth subscript p stands for parity (0 for real part, and 1 for imaginary). Each bracketed term corresponds to a 1D filter (acting either on the x or y directions). For the horizontal and vertical orientations ($\theta_0 = 0^0$ and 90^0) we need two 1D Gabor filters (for $p = 0$ and $p = 1$), tuned to the frequency $f_0 = (1/2)f_N$, plus a 1-D Gaussian filter. For the two main diagonals ($\theta_0 = 45^0$ and 135^0), we again need 2 filters for the two parities, but now tuned to the frequency $\left(\sqrt{2}/2\right)f_0$.

The theoretical DC (zero frequency) response of the real part of our 2-D Gabor filters is close to, but not exactly zero. Thus, we have imposed a zero DC response on them. We have performed a constrained least squares minimization of the error of the frequency response, using the Moore-Penrose generalized inverse [23] to design the four one-dimensional Gabor masks. We have chosen a size of 11 elements for the masks, that provides high fidelity approximating the original filters while maintaining low computational cost. Table 1 contains the resulting values of the integer arithmetic filters used, showing also their SNR. The errors are larger for the even filters due to the additional zero DC constraint. The SNR of the equivalent 2-D frequency responses respect to the original 2-D Gabor filters is 25.1 dB for the diagonal filters, 27.3 dB for the even horizontal and vertical and 31.8 dB for the odd horizontal and vertical

############# Insert Table 1 about here #############

5

High Pass Residual

As we said in the Introduction, there are many applications where the outputs of the set of Gabor filters, plus the LPR, suffices to obtain a good representation of the input image. However, if we want to achieve a complete representation of the image, we need to include an HPR. The ideal residual filter, showed in Eq. 3, compensates for both the loss of the very high frequency components and the amplitude distortion. Strictly speaking, such a filter is a residual but it is not purely high-pass. In this Section we analyze two possibilities for approximating the ideal residual: 1) as a single filter in the spatial domain, using a reduced mask, and 2) applying an approximated, highly efficient method, which provides a purely high-pass HPR.

Following the first approach, we have implemented the impulse response in Eq. 3 as a 11x11 convolution mask. We have used the same least-square error method for the filter design, extended for two dimensions, and including the zero DC response constraint. The resulting mask is shown in Table 2. Unfortunately, this filter is not separable. Nevertheless, we can save much computing time if we take advantage of its four-quadrant symmetry. In order to be more consistent with our 1D filtering implementation, we have also approximated this filter as the sum of four separable masks (see the Appendix). This can be useful when using special computer architectures based on one-dimensional convolvers.

############## Insert Table 2 about here ##############

There is an exact (and rather trivial) solution for the HPR design problem, consisting of subtracting all the Gabor channels and the LPR, conveniently stretched to the original size, from the original image. However, this procedure requires the reconstruction of the input image from the channels and for this reason it has not been considered here. A more practical approach has been followed in this work, that consist of subtracting from the original only the highest frequency channels, plus the low-pass filtered version of the image obtained in the second level of the pyramidal implementation. This method is highly efficient, since it is only required to make a subtraction and produces a pure HPR (in the sense that it just compensates for the loss of the high frequencies). We have scaled the four first-level Gabor channel responses before subtracting them to avoid the possibility of a phase inversion in the resulting HPR (a scale factor of 0.75 is suitable for this purpose).

Figure 3 shows the Gabor decomposition of a 256x256 test image, including the HPR (computed by the subtraction method), the LPR and the original image. The mid gray level represents the zero response (except for the original image and the LPR, which keep the original gray scale), lighter gray levels mean positive response and darker ones negative. It is clear the different nature of the residuals (selective to frequency but not to orientation) with respect to the Gabor channels (selective to orientation and frequency).

############## Insert Figure 3 about here ##############

4.- IMAGE RECONSTRUCTION

One of the main features of this scheme (different from most approaches) is that the reconstruction is very simple and direct. The pyramidal implementation produces a redundant but robust representation. We obtain a good reconstruction by simply adding together all the even Gabor channels (previously interpolated to their original sizes, as explained in [8]) plus the residuals, in the same fashion as in other similar multiresolution codes [20, 21].

Equalization

A fixed scale factor is independently applied to each Gabor channel before their addition, trying to compensate both for the non flat response of the sum of the Gabor filters and for the effect of the low pass filter used in the pyramidal implementation.

We have again applied a constrained least squares optimization to obtain the weights of each channel, including both LPR and HPR that have been forced to maintain their theoretical gains of 1. Since we have argued that in many situations the HPR will not be necessary [8], we have computed the set of equalization weights for the cases of including or not the HPR. In the second case we have only considered the region of the Fourier domain well covered by the 4x4 Gabor channels and the LPR (a circle with a radius $0.656\,f_N$, that corresponds to the 3 dB bandwidth of the response of the whole system). In this way we avoid trying to compensate for the loss of the very high frequencies, which are not properly covered by the set of Gabor channels. Table 3 contains the resulting optimal weights.

############# **Insert Table 3 about here** #############

Objective and subjective quality of the reconstruction

Figure 4 shows the reconstruction results obtained with and without the HPR (computed using the subtraction method), as well as the reconstruction errors. If we do not include the HPR (panel a), the quality of the reconstruction is very good, although there is a slight low-pass filtering effect. If we do include the HPR (panel b), the reconstruction is visually indistinguishable from the original.

############# **Insert Figure 4 about here** #############

For an overall objective evaluation of the fidelity of the reconstruction, we have used a set of seven test images (namely, Einstein and Lena portraits, a moon picture, an x-ray view of human bones, a straw texture sample, a spatially variant sinusoid and a white noise sample). Table 4 summarizes the signal-to-noise ratios (in dB) of the reconstruction. The three rows correspond, respectively, to a natural test image (Einstein), white noise, and the average result for the seven test images. The three columns correspond to the three considered options (without HPR, with HPR computed using the subtraction method and with HPR computed using the 11x11 mask).

############# **Insert Table 4 about here** #############

A remarkable fact is the large improvement in the SNR for the white noise (from 2.3 dB to 24 or 29 dB, depending on the HPR implementation). The low SNR

obtained in the case when we do not consider the HPR can be expected because we are discarding a relevant portion of the spectrum. However, the improvement is more modest (4 dB) for the Einstein portrait. The reason is that white noise is not a typical example of a natural image, since natural scenes tend to show an exponential decay in their power spectrum [19]. The Einstein portrait is a much more representative sample of natural images. The reconstruction improves about 6 dB in average when the HPR channel is included.

If we compare the results obtained with the two implementations of the HPR, it turns out that they are completely equivalent when applied to natural images, whereas the 11x11 filter mask performs better (by 5 dB) when applied to white noise. However, this solution represents a minor improvement, having a very high relative cost (almost doubling the original cost).

5.- COMPUTATIONAL ADVANTAGES OF THE SPATIAL DOMAIN IMPLEMENTATION

For the spatial domain implementation, we must know the number of filtering operations needed to obtain the 4 orientations and 2 parities of a level. Figure 5 shows a block diagram (derived from Eq. 4) of these filtering operations. We need to perform twelve 1D convolutions to obtain a complete level, whereas we only need eight to obtain one parity (because some convolutions are common to both parities). Therefore, the total number of real operations (additions and multiplications) required to obtain the 4x4 even channels is:

$$\left[4(3n-1)+(3n'-1)\right]\sum_{k=0}^{3}\left(\frac{N}{k}\right)^2 + 6N^2 \cong \frac{4}{3}\left[4(3n-1)+(3n'-1)\right]N^2 + 6N^2, \qquad (5)$$

where NxN are the image dimensions, n is the dimension of the Gabor filter masks and n' is the dimension of the low-pass filter mask. The last term accounts for the HPR computation (using the subtraction method).

############# Insert Figure 5 about here #############

For the Fourier domain implementation, and following the complexity analysis made by Watson [20], the total number of real operations required to obtain the 4x4 even channels is:

$$5N^2\log_2 N + 4\sum_{k=0}^{3}\left(\frac{N}{k}\right)^2\left[1+5(\log_2 N - k)\right] + 5N^2\log_2 N + N^2, \qquad (6)$$

where the two last terms account for the HPR computation.

The plot in Figure 6 shows a comparison between the number of operations per pixel required for the frequency and space domain implementations (taking $n=11$ and $n'=5$), and for several image sizes. The number of operations grows logarithmically with the image size for the Fourier implementation, and it is always higher than the number of operations for the spatial implementation (which remains constant).

8

############### **Insert Figure 6 about here** ###############

Another important issue is the amount of memory needed to store the filter responses. The FFT implementation needs to keep in memory the frequency responses of the highest frequency filters and the HPR, for the largest image size required. This imposes a serious limitation on this option. On the other hand, our implementation only has to keep in memory a few coefficients corresponding to the 1D convolution masks, without limiting the size or shape of the input image.

6.- SUMMARY

We have proposed a pyramidal multiscale Gabor wavelet transform that permits a very efficient implementation in the spatial domain, which is faster than conventional FFT implementations. Contrary to most wavelet schemes, our moderately redundant approach permits a very fast reconstruction of the image, by simply adding the set of weighted channels. The weights are fixed and have been designed to minimize the reconstruction error.

In the original scheme [8] there was a low-pass filtering effect in the reconstruction, due to the problem of the loss of information produced when sampling a squared domain using a log-polar grid. This has been solved by incorporating a high-pass residual. We have implemented and analyzed two possible solutions in the spatial domain. The fastest one consists of subtracting from the original image its low-pass filtered version (obtained in the first level of the pyramid, before down sampling), together with the highest frequency Gabor channels. The other option (much more computationally expensive) consists of applying a 11x11 filter mask. Although this second option is more consistent with our spatial domain scheme, the small improvement in fidelity with respect to the former one hardly justifies the large increment in computational cost, at least when dealing with natural images. However, this option could be interesting when working with special architectures that perform high speed 1D convolutions (see Appendix).

In conclusion, we have developed an optimized spatial domain implementation of a multipurpose, multiscale Gabor scheme, taking advantage of the special properties (some of them unique) of the Gabor functions. We have achieved a very low computational cost and a high fidelity both in the filter design and in the reconstruction.

APPENDIX

One of the proposed options for implementing the HPR is as a non separable 11x11 convolution mask. We have decomposed this mask as the sum of separable ones, to be consistent with our implementation based on 1D filters. We have applied the Singular Value Decomposition (SVD) [24] to the two dimensional mask in Table 2, obtaining the following six non-zero singular values: 182.3, 69.4, 22.7, 15.3, 2.5 and 0.8. The two last SV are negligible compared to the others. Therefore, we can accurately approximate the original 11x11 mask as the following sum of only 4 separable products,

$$hpr \approx \frac{1}{256^2}(\mathbf{AA}^T - \mathbf{BB}^T + \mathbf{CC}^T - \mathbf{DD}^T), \qquad (7)$$

where the resulting 1D filters, **A**, **B**, **C** and **D** include the scaling factor of their corresponding singular value, and are showed in Table 5. We have verified, using the white noise image, that there is no noticeable difference between using the 11x11 mask vs. this sum of 4 separable filters.

############# Insert Table 5 about here #############

Acknowledgments

REFERENCES

[1]   R. Navarro, A. Tabernero, and G. Cristóbal, "Image Representation with Gabor Wavelets and Its Applications," in *Advances in Imaging and Electron Physics*, vol. 97, P. W. Hawkes, Academic Press, San Diego CA, pp. 1-84 (1996).

[2]   J. G. Daugman, "Complete discrete 2D Gabor transform by neural networks for image analysis and compression," *IEEE Trans. on Acoust., Speech, and Signal Proc.* 36, 1169-1179 (1988).

[3]   G. Cristobal and R. Navarro, "Space and frequency variant image enhancement based on a Gabor representation," *Patt. Recog. Lett.*, 15, 273-7 (1994).

[4]   M. Turner, "Texture discrimination by Gabor functions," *Biol. Cyber.*, 55, 71-82 (1986).

[5]   O. Nestares, R. Navarro, J. Portilla, and A. Tabernero, "Automatic computation of the area irradiated by ultrashort laser pulses in Sb materials through texture segmentation of TEM images," *Ultramicroscopy*, (in press).

[6]   J. Portilla, R. Navarro, O. Nestares, and A. Tabernero, "Texture synthesis-by-analysis based on a multiscale early-vision model," *Optical Engineering*, 35, 2403-2417 (1996).

[7]   D. J. Heeger, "Model for the extraction of image flow," *J. Opt. Soc. Am.*, A 4, 1455-1471 (1987).

[8]   R. Navarro and A. Tabernero, "Gaussian wavelet transform: two alternative fast implementations for images," *Multidimensional System and Signal Processing*, 2, 421-436 (1991).

[9]   A. C. Bovik, N. Gopal, T. Emmoth, and A. Restrepo, "Localized measurement of emergent image frequencies by Gabor wavelets," *IEEE Trans. on Inf. Theory*, 38,2, 691-712 (1992).

[10]   M. Porat and Y. Y. Zeevi, "The generalized Gabor scheme for image representation in biological and machine vision," *IEEE Trans. on Patt. Anal. and Mach. Intell.*, 10, 452-468 (1988).

[11]   J. G. Daugman, "Spatial visual channels in the Fourier plane," *Vis. Res.*, 24, 891-910 (1984).

[12]   S. Marcelja, "Mathematical description of the response of simple cortical cells," *J. Opt. Soc. Am.*, A 70, 1297-1300 (1980).

[13]   J. G. Daugman, "Uncertainty relation for resolution in space, spatial--frequency and oriented optimized by two--dimensional visual cortical filters," *J. Opt. Soc. Am. A*, 2, 1160-1169 (1985).

[14]   M. J. Bastiaans, "A sampling theorem for the complex spectrogram, and Gabor's expansion of a signal in Gaussian elementary signals," *Opt. Eng.*, 20, 594-598 (1981).

[15]   R. N. Braithwaite and M. P. Beddoes, "Iterative methods for solving the Gabor expansion: considerations of convergence," *IEEE Trans. on Image Processing*, 1,2, 243-234 (1992).

[16]   T. Ebrahimi and M. Kunt, "Image compression by Gabor expansion," *Opt. Eng.*, 30, 873-880 (1991).

[17]   R. A. Young, "The Gaussian derivative model for spatial vision: I. Retinal mechanisms," *Spatial Vision*, 2, 273-293 (1987).

[18]   E. P. Simoncelli, W. T. Freeman, E. H. Adelson, and D. J. Heeger, "Shiftable multi-scale transforms," *IEEE Trans. on Inf. Theory*, 38, 587-607 (1992).

[19]   D. J. Field, "Relation between the statistics of natural images and the response properties of cortical cells," *J. Opt. Soc. Am. A*, 4, 2379-2394 (1987).

[20]   A. B. Watson, "The cortex transform: rapid computation of simulated neural images," *Comp. Vis., Graph., and Image Proc.*, 39, 311-327 (1987).

[21]   P. L. Burt and E. H. Adelson, "The Laplacian pyramid as a compact image code," *IEEE Trans. on Comm.*, 31, 532-540 (1983).

[22]   W. K. Pratt, *Digital Image Processing*, Wiley, New York(1991).

[23]   R. Gonzalez and P. Wintz, *Digital Image Processing*, Addison-Wesley, Reading, Mass.(1987).

[24]   A. K. Jain, *Fundamentals of Digital Image Processing*, Prentice-Hall, Inc, New Jersey(1989).

TABLES

Table 1. Filter masks used in the implementation of the Gabor wavelet transform. Values are given for integer arithmetic; the filter responses must be normalized by 256 (except for the last mask, for which the normalization factor is 16). Fidelity of the design is given as signal-to-noise ratio in the last column.

Filter	Mask	SNR
H-V (main directions) Even	0 17 0 -62 0 90 0 -62 0 17 0	28.9 dB
H-V (main directions) Odd	-8 0 37 0 -82 0 82 0 -37 0 8	40.2 dB
Diagonal Even	3 -5 -27 -27 25 62 25 -27 -27 -5 3	25.7 dB
Diagonal Odd	4 13 5 -34 -52 0 52 34 -5 -13 -4	35.6 dB
Pure Gaussian	4 9 19 31 41 45 41 31 19 9 4	32.4 dB
Low-Pass (cubic B-spline)	1 4 6 4 1	

Table 2.- Coefficients for the 11x11 HPR mask (for integer arithmetic and scaled by 256).

0	1	2	3	0	-3	0	3	2	1	0
1	0	-1	0	-3	-6	-3	0	-1	0	1
2	-1	-7	-3	3	4	3	-3	-7	-1	2
3	0	-3	7	11	7	11	7	-3	0	3
0	-3	3	11	-19	-49	-19	11	3	-3	0
-3	-6	4	7	-49	160	-49	7	4	-6	-3
0	-3	3	11	-19	-49	-19	11	3	-3	0
3	0	-3	7	11	7	11	7	-3	0	3
2	-1	-7	-3	3	4	3	-3	-7	-1	2
1	0	-1	0	-3	-6	-3	0	-1	0	1
0	1	2	3	0	-3	0	3	2	1	0

Table 3.- Optimal gains of the Gabor channels for equalization of the frequency response, including/not including the HPR.

	0^0	45^0	90^0	135^0
f_1	1.52 / 1.53	1.27 / 1.24	1.52 / 1.53	1.27 / 1.24
f_2	1.77 / 1.88	1.48 / 1.64	1.77 / 1.88	1.48 / 1.64
f_3	1.56 / 1.68	1.45 / 1.43	1.56 / 1.68	1.45 / 1.43
f_4	1.06 / 1.30	1.10 / 1.25	1.06 / 1.30	1.10 / 1.25

Table 4.- Signal-to-noise ratios obtained in the reconstruction of 7 test images, including Einstein's portrait and white noise. We compare the reconstruction results without considering the HPR (first column) with the two alternative implementations of the HPR.

	NO HPR	HPR (subtraction)	HPR (11x11 filter)
Einstein	24.7 dB	28.4 dB	27.4 dB
White Noise	2.3 dB	24.0 dB	29.2 dB
Mean (7 images)	**21.0 dB**	**26.7 dB**	**26.8 dB**

Table 5.- Resulting masks for the separable implementation of the HPR (for integer arithmetic and scaled by 256).

A	3	5	-3	-2	45	-206	45	-2	-3	5	3
B	-5	-10	15	27	-84	-39	-84	27	15	-10	-5
C	12	-4	-4	50	14	5	14	50	-4	-4	12
D	-13	11	40	5	7	2	7	5	40	11	-13

FIGURE CAPTIONS

Fig. 1.- Real part of the 4x4 Gabor functions: (a) Functions in the spatial domain; (b) Coverage of the Fourier domain by their corresponding frequency channels.

Fig. 2.- Spatial impulse response of the ideal HPR, after Eq. 3.

Fig. 3.- Gabor decomposition of a test image. The upper pyramid represents the even (real) channels, including both the HPR (upper left), and the LPR (small square in the center left). The lower pyramid displays the odd (imaginary) part.

Fig. 4.- Image reconstruction results (upper) and their corresponding absolute errors (lower). The errors have been scaled (by a factor of 3) to facilitate visualization: (a) Reconstruction without including HPR; (b) Reconstruction with HPR.

Fig. 5.- Block diagram for the filtering operations used to obtain the 4 orientations and the 2 parities of a level. Each block represents a 1-D filtering operation in the direction specified (X or Y): (a) Vertical and horizontal directions; (b) Diagonal directions

Fig. 6.- Number of real operations per pixel required for frequency (via FFT) and spatial domain (taking $n=11$ and $n'=5$) implementations, for an image size of NxN.