Matching pursuit of images

François Bergeaud
and Stéphane Mallat

Abstract. A crucial problem in image analysis is to construct efficient low-level representations of an image, providing precise characterization of features wich compose it, such as edges and texture components.

An image usually contains very different types of features, which have been successfully modelled by the very redundant family of 2D Gabor oriented wavelets, describing the local properties of the image: localization, scale, preferred orientation, amplitude and phase of the discontinuity.

However, this model generates representations of very large size. Instead of decomposing a given image over this whole set of Gabor functions, we use an adaptative algorithm (called Matching Pursuit) to select the Gabor elements wich approximate at best the image, corresponding to the main features of the image .

This produces compact representation in terms of few features that reveal the local image properties. Results proved that the elements are precisely localized on the edges of the images, and give a local decomposition as linear combinations of "textons" in the textured regions.

We introduce a fast algorithm to compute the Matching Pursuit decomposition for images with a complexity of $\mathcal{O}(N \log^2 N)$ per iteration for an image of N^2 pixels.

Matching Pursuit of Images
François Bergeaud and Stéphane Mallat, pp. 1–16.

§1 Introduction

The complexity of image structures including different types of textures and edges requires flexible image representations. Although an image is entirely characterized by its decomposition in a basis, any such basis is not rich enough to represent efficiently all potentially interesting low-level structures. Some image components are diffused across many bases elements and are then difficult to analyze from the basis representation. This is like trying to express oneself in a language including a small dictionary. Non available words must be replaced by long awkward sentences. To provide explicit information on important local properties, the image is represented as a sum of waveforms selected from an extremely redundant dictionary of oriented Gabor functions. As opposed to previous approaches, we do not decompose the image over the whole dictionary, but like in a sentence formation, we select the most appropriate Gabor waveforms to represent the image. Instead of increasing the representation by a large factor as in typical multiscale Gabor representations ([14]), the adaptive choice of dictionary vectors defines a compact representation that takes advantage of the flexibility offered by the dictionary redundancy.

There is an infinite number of ways to decompose an image over a redundant dictionary of waveforms. The selection of appropriate waveforms to construct the image representation is obtained by constructing efficient image approximations from few dictionary vectors. The optimization of the approximation is not intended for data compression but as a criteria for feature selection. If most of the image is recovered as a sum of few dictionary vectors, these vectors must closely match the local image properties. One can however prove that finding optimal approximations in redundant dictionaries is an NP complete problem. The redundancy opens a combinatorial explosion. This explosion is avoided by the matching pursuit algorithm that uses a non optimal greedy strategy to select each dictionary elements. For a dictionary of Gabor functions, the greedy optimization of the image approximation leads to an efficient image representation where each Gabor waveform reflects the orientation, scale and phase of local image variations. For textures, the selected Gabor elements can be interpreted as textons where as along edges, the multiscale properties of these Gabor elements reflect the edge properties. When the image is translated or rotated, the selected Gabor elements are translated and rotated. A fast implementation of this algorithm and numerical examples are presented.

§2 The 2D Gabor wavelet dictionary

Image decompositions in families of Gabor functions characterize the local scale, orientation and phase of the image variations. Gabor functions are constructed from a window $b(x,y)$, modulated by sinusoidal waves of fixed frequency ω_0 that propagate along different direction θ with different phases ϕ

$$b_{\theta,\phi}(x,y) = b(x,y)\cos(\omega_0(x\cos\theta + y\sin\theta) + \phi). \qquad (2.1)$$

Each of these modulated windows can be interpreted as wavelets having different orientation selectivities. Figure (1) shows the surface of such a 2D Gabor wavelet for $\phi = \theta = 0$. The window $b(x,y)$ is not chosen to be a Gaussian but is a compactly supported box spline that is adjusted so that the average of $b_{\theta,\phi}(x,y)$ is zero for all orientations and phases. The phase modifies the profile of the window oscillations. When $\phi = 0$, the oscillations are symmetrical and the window can be interpreted as the second order derivative of a two-dimensional window, along the direction θ. When $\phi = \pi/2$, the oscillations are antisymmetrical and the corresponding wavelet rather behaves as a first order derivative along the direction θ. The figure (2) shows the cross-section of an element of the dictionary for different values of the phase ($\phi = -\frac{\pi}{4}$, $\phi = 0$, $\phi = \frac{\pi}{4}$, $\phi = \frac{\pi}{2}$).

These oriented wavelets are then scaled by s and translated to define a whole family of Gabor wavelets $\{g_\gamma\}_{\gamma \in \Gamma}$ with:

$$g_\gamma(x,y) = \frac{1}{s} b_{\theta,\phi}\left(\frac{x-u}{s}, \frac{y-v}{s}\right) \tag{2.2}$$

where the multi-index parameter $\gamma = (\theta, \phi, s, u, v)$ carries the orientation, phase,scale and position of the corresponding Gabor function.

The Gabor transform of an image $f(x,y)$ is defined by the inner product

$$Gf(\gamma) = < f, g_\gamma > = \iint_{\mathrm{R}^2} f(x,y) g_\gamma(x,y)\, dx\, dy.$$

The orientation parameter θ allows to match the different orientation of image structures. The scaling factor s allows to zoom into singularities at fine scales but also to recover large scale image variations. The phase is a convenient parameter to modify the number of vanishing moments of the wavelet. When $\phi = 0$, the Gabor transform provides a second order partial derivative along θ of the image that is smoothed by a dilated window. When $\phi = \pi/2$, the Gabor transform can be interpreted as a first order partial derivative and thus responds particularly to discontinuities and edges.

The Fourier transform of a Gabor function is a waveform whose energy is well concentrated in the Fourier plane. In numerical computations, the scale is restricted to powers of two $\{2^j\}_{j \in \mathrm{Z}}$ and the angles are discretized. The Gabor dictionary used in this paper includes 8 orientations (cf. figure (3)). To define a complete representation, we guaranty that the whole Fourier plane is covered by dilations of the 8 elementary Gabor wavelets.

Besides engineering and mathematical motivations, families of Gabor waveforms have been used as models for the receptive fields of simple cells. The representation of the visual information is described as moving from the retina to the primary visual cortex: at the retina level, the receptive fields are homogenous and provide a broad bandwidth for the visual information received from the photoreceptors. In the cortex, the receptive fields are narrow-band and direction-oriented. The cortex has multiple representations, generated by distinct neural populations.

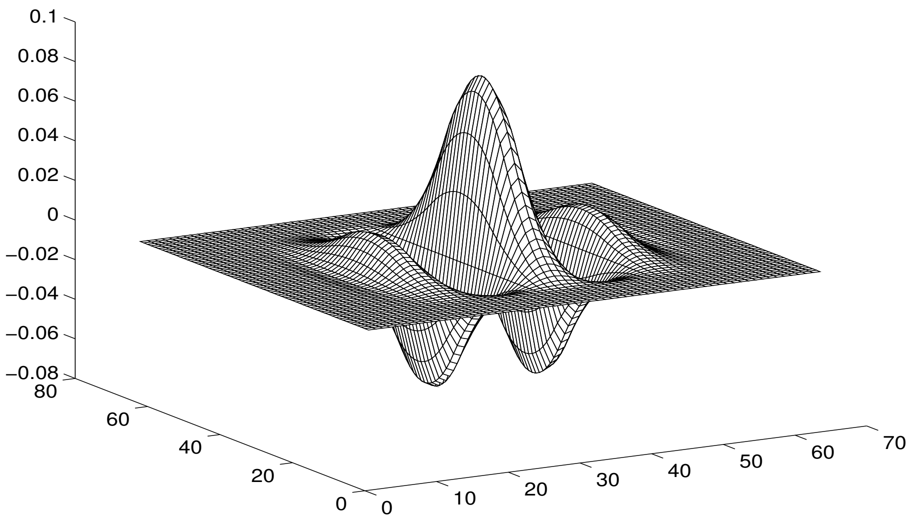

Figure 1. The surface of the 2D Gabor function

The properties of simple cells of the V1 area of the cortex suggests some requirements for the image representation: they perform a linear spatial summation of their inputs. They are local in space and frequency and cover the entire spectral visual field. They are divided among several types, each of which being selective to the orientation (with a maximum range of 45°, which suggests a minimum of 4 orientations), and having a limited frequency bandpass (1.5 octave), which requires self similarity among the different V1 cells class. The receptive fields profile of simple cells, which appears to be the descriptors of the features have been shown to be strikingly similar to the two-dimensional Gabor functions ([13]).

The decomposition of images in a Gabor dictionary defines a very redundant representation. For an image of 512 by 512 pixels, 8 orientations, 6 octaves and a two phases ($\phi = 0$ and $\phi = \frac{\pi}{2}$) representation would correspond to 96 images of 512 by 512 pixels. These images could be subsampled but we then lose the translation invariance of the representation. Instead of decomposing the image over the whole dictionary, we select specific Gabor waveforms that provide an efficient image approximation.

§3 Matching Pursuit

We consider the general problem of decomposing a signal f over a dictionary of unit vectors $\{g_\gamma\}_{\gamma\in\Gamma}$ whose linear combinations are dense in the signal space $\mathcal{H}$.

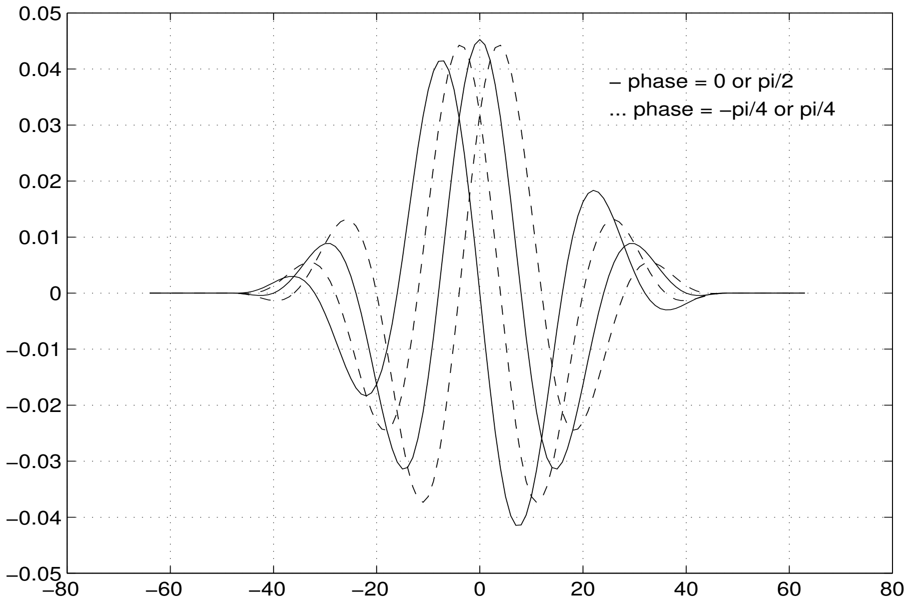

Figure 2. Cross-section of a 2D Gabor function for $\phi = -\frac{\pi}{4}$, 0, $\frac{\pi}{4}$, $\frac{\pi}{2}$

The smallest possible dictionary is a basis of $\mathcal{H}$; general dictionaries are redundant families of vectors. When the dictionary is redundant, unlike the case of a basis, we have some degree of freedom in choosing a signal's particular representation. This freedom allows us to choose few dictionary vectors, whose linear combinations approximate efficiently the signal. The chosen vectors highlights the predominant signal features. For any fixed approximation error ϵ, when the dictionary is redundant, we can show that finding the minimum number of dictionary elements that approximates the image with an error smaller than ϵ is an NP hard problem.

Because of the difficulty of finding optimal solutions, we use a greedy matching pursuit algorithm that has previously been tested on a one-dimensional signal.The matching pursuit ([10]) uses a greedy strategy that computes a good suboptimal approximation. It successively approximates a signal f with orthogonal projections onto dictionary elements. The first step is to approximate f by projecting it on a vector $g_{\gamma_0} \in \mathcal{D}$:

$$f = < f, g_{\gamma_0} > g_{\gamma_0} + Rf, \tag{3.1}$$

Since the residue Rf is orthogonal to g_{γ_0},

$$\|f\|^2 = |< f, g_{\gamma_0} >|^2 + \|Rf\|^2. \tag{3.2}$$

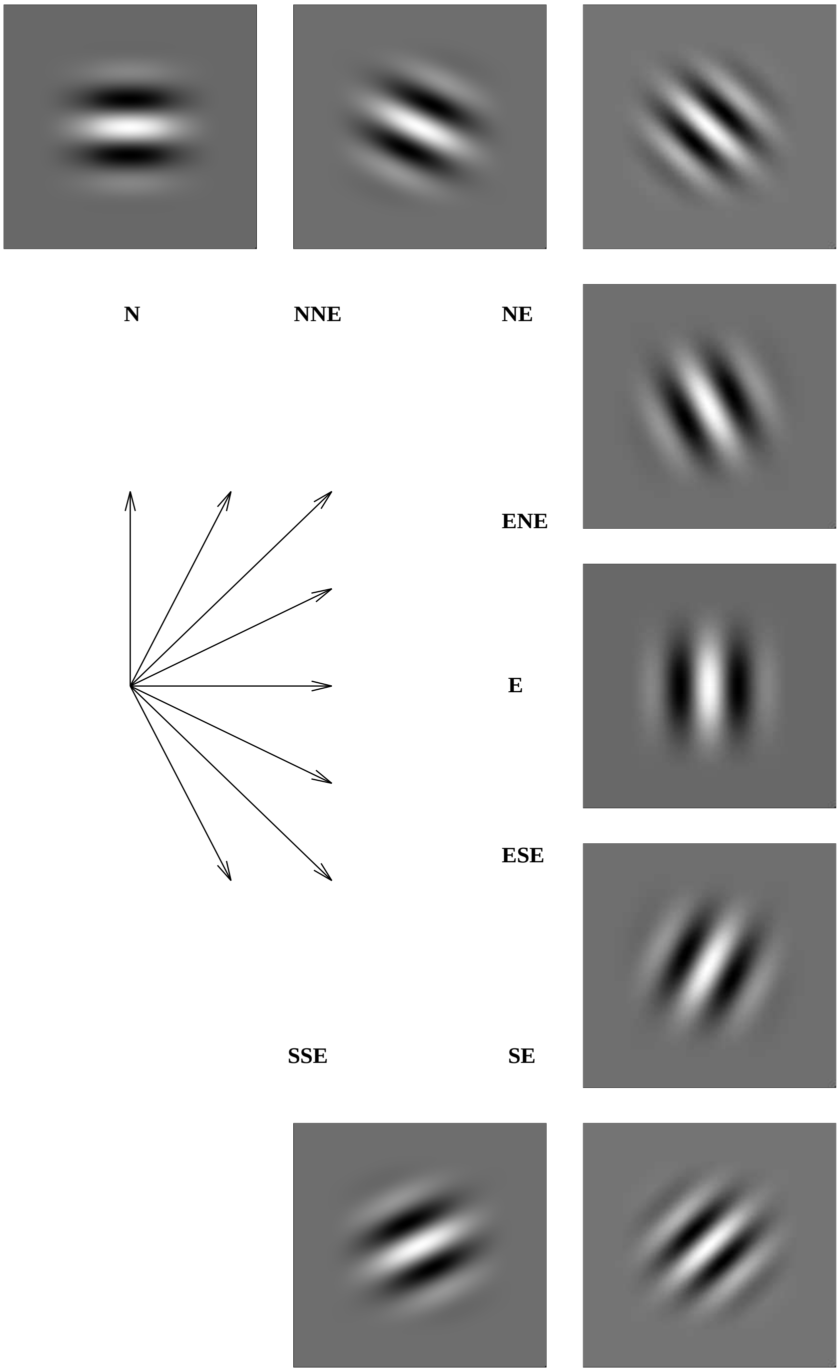

Figure 3. Discretization of the angle θ

We minimize $\|Rf\|$ by choosing g_{γ_0} which maximizes $|< f, g_\gamma >|$. We choose g_{γ_0} such that

$$|< f, g_{\gamma_0} >| = \sup_{\gamma\in\Gamma} |< f, g_\gamma >|, \tag{3.3}$$

The pursuit iterates this procedure by sub-decomposing the residue. Let $R^0 f = f$. Suppose that we have already computed the residue $R^k f$.

We choose $g_{\gamma_k} \in \mathcal{D}$ such that:

$$|< R^k f, g_{\gamma_k} >| = \sup_{\gamma\in\Gamma} |< R^k f, g_\gamma >|. \tag{3.4}$$

and project $R^k f$ on g_{γ_k}

$$R^{k+1} f = R^k f - < R^n k f, g_{\gamma_k} > g_{\gamma_k}. \tag{3.5}$$

which defines the residue at the order k+1. The orthogonality of $R^{k+1}f$ and g_{γ_k} implies

$$\|R^{k+1} f\|^2 = \|R^k f\|^2 - |< R^k f, g_{\gamma_k} >|^2. \tag{3.6}$$

By summing (3.5) for k between 0 and $n-1$, we obtain

$$f = \sum_{k=0}^{n-1} < R^k f, g_{\gamma_k} > g_{\gamma_k} + R^n f. \tag{3.7}$$

Similarly, summing (3.6) for k between 0 and $n-1$ yields

$$\|f\|^2 = \sum_{n=0}^{n-1} |< R^k f, g_{\gamma_k} >|^2 + \|R^n f\|^2. \tag{3.8}$$

The residue $R^n f$ is the approximation error of f after choosing n vectors in the dictionary and the energy of this error is given by (3.8).

In infinite dimensional spaces, the convergence of the error to zero is shown ([10]) to be a consequence of a theorem proved by Jones ([7]):

$$\lim_{m\to+\infty} \|R^m f\| = 0. \tag{3.9}$$

Hence

$$f = \sum_{n=0}^{+\infty} < R^n f, g_{\gamma_n} > g_{\gamma_n}, \tag{3.10}$$

and we obtain an energy conservation

$$\|f\|^2 = \sum_{n=0}^{+\infty} |< R^n f, g_{\gamma_n} >|^2. \tag{3.11}$$

In finite dimensional signal spaces, the convergence is proved to be exponential ([10]).

§4 Fast calculation

Despite the apparent brute force strategy of a matching pursuit, this algorithm can be implemented in an efficient way that requires on average $\mathcal{O}(N\,\log^2 N)$ operations per iteration in a Gabor dictionary, for an image of N^2 pixels.

The fast algorithm is based on three major points: as the search for the maximal coefficient (scalar product of the residue with a dictionary vector) is very time-consuming, we first make all the computations on a much smaller dictionary $\mathcal{D}' \subset \mathcal{D}$. It is constructed such that the maximal coefficient in $\mathcal{D}'$ gives a good estimation of the maximal coefficient in the whole dictionary. We then use a local Newton search of complexity $\mathcal{O}(1)$ to find the best local coefficient of $\mathcal{D}$ in the neighborhood of the element selected in $\mathcal{D}'$.

As we want to find the maximal coefficient in $\mathcal{D}'$ in an quick and efficient way, we sort and store the scalar products of the current residue with the elements of $\mathcal{D}'$ in a Hash table. The maximal coefficient is always at the top of the table, so that we don't need to search for it.

The coarse dictionary $\mathcal{D}' = (g_{\gamma'})_{\gamma' \in \Gamma'}$ is obtained by subsampling the position of the elements of the dictionary by a factor of 4 at each scale. It is nearly a tight frame of $\mathcal{H}$ and thus complete, because we only remove the excess of redundancy from the initial dictionary. The size of such a coarse dictionary goes down to $\frac{64}{3}\,N^2$ (instead of $16\,N^2\,\log N$ for the whole dictionary $\mathcal{D}$).

Once we have an estimation of the best vector g_{γ_n} such that:

$$|< R^n f, g_{\gamma_n} >| = \sup_{\gamma \in \Gamma} |< R^n f, g_\gamma >| \tag{4.1}$$

we compute the inner product of the new residue $R^{n+1} f$ with any $g_{\gamma'} \in \mathcal{D}'$, with a linear updating formula derived from equation (3.5)

$$< R^{n+1} f, g_{\gamma'} > = < R^n f, g_{\gamma'} > - < R^n f, g_{\gamma_n} > < g_{\gamma_n}, g_{\gamma'} > . \tag{4.2}$$

This updating equation can be interpreted as an inhibition of $< R^n f, g_\gamma >$ by the cross correlation of g_{γ_n} and g_γ. The only coefficients $< R^n f, g_{\gamma'} >$ of the Hash table wich are modified are such that $< g_{\gamma_n}, g_{\gamma'} >$ is non zero. As we use compact supported Gabor wavelets, this set is indeed very sparse, and the updating formula only apply to few coefficients ($\mathcal{O}(N\,\log^2 N)$ in the average), compared to the size of the table ($\mathcal{O}(N^2)$) . We store the set of all non null inner products $< g_\gamma, g_{\gamma'} >$. Since the Gabor Wavelets used in these numerical experiments are separable, we only need to compute and store the set of inner products of 1D Gabor functions, thus requiring an order of magnitude less memory.

§5 Results

The matching pursuit algorithm applied to a Gabor dictionary selects iteratively the Gabor waveforms, also called atoms, whose scales, phases, orientations and positions best match the local image variations.

The distribution of Gabor functions at the scales 2^j for $1 \leq j \leq 4$ is shown in fig. (4, 5). For the lady image (256×256), light atoms correspond to high amplitude inner products. As clearly shown by these images, the fine scale atoms are distributed along the edges and the texture regions and their orientation indicate the local orientation of the image transitions.

In order to display the edge information (localization, orientation, scale, amplitude), we adopt the following convention: each selected Gabor vector g_γ for $\gamma_n = (\theta, \phi, 2^j, u, v)$ is symbolized by an elongated Gaussian function of width proportional to the scale 2^j, centered at (u, v), of orientation θ. The mean gray level of each symbol is proportional to $| < R^n f, g_{\tilde{\gamma_n}} > |$.

Figure (6) displays the reconstruction of the Lena image with the 2500 and 5000 first Gabor waveforms selected by the algorithm. The image reconstruction is simply obtained with the truncated sum of the infinite serie decomposition (3.10). With 2500 atoms, we already recover a good image quality and with 5000 atoms the reconstructed image has no visual degradation. The latter image is constructed without using any compression method, and achieves a 1.5 bit/pixel ratio with a perfect visual quality. This shows clearly that, while specifying precisely the important information such as the local scale and orientation, this representation is very compact and can be used as the input to a high level processing. Moreover, the degradation of the image quality is very small when using 2500 atoms instead of 5000 atoms, which suggests that the most relevant features of the image lie in a very small set of atoms.

The fig. 7 illustrate the texture discrimination properties of the representation: a straw texture image is inserted into a paper texture image. The paper texture has no orientation specificity. On the other hand, the straw texture has horizontal and vertical structures. At fine scales, most structures are vertical because the horizontal variations are relatively smooth. At the intermediate scale 2^2 we clearly see the horizontal and vertical image structures. At the larger scale 2^3, the vertical structures dominate again because the vertical straw are of larger sizes.

The symbolic representations at scale 2^1, 2^2, 2^3, 2^4, show the distinct behaviour of the two textures relative to the scale: although the paper texture shows a scale-invariant uniform distribution, the straw texture representation points out the vertical features of the straw at finer and larger scale, and the horizontal and vertical structures of the straw at medium scale. Moreover, at scale 2^1, we clearly distinguish the texture edge as the boundary between the two textures.

As noticed by Turner in ([14]), the time-frequency analysis performed using Gabor functions corresponds to a tiling of the 4-dimentional space (2D for

Figure 4. Symbolic representation of the Gabor vectors: each vector is symbolized by an elongated Gaussian of width proportional to the scale 2^j, oriented along the modulation direction, of gray level proportional to the amplitude $|< R^n f, g_{\gamma_n} >|$. Light atoms correspond to high amplitude. The left figure corresponds to the scale 2^1 and the right figure to the scale 2^2.

spatial position, 2D for the frequency position) called "information hyperspace". Each element of the dictionary occupies a particular "volume" or cell in this space, and combine information in space and frequency related to the structural description of the texture: the response to the spatially large Gabor filters and small spatial frequency extent give a precise information on the periodicity of the texture. On the other hand, the response to the spatially small filters and large spatially frequency extent distinguish texture elements and provide a mean to characterize the density of elements. The intermediate Gabor filters allow simultaneous measurements of textons and their distribution.

The fundamental idea, introduced in previous work by Julesz ([9]) and Beck ([1]), and stated by Malik and Perona in ([11]) is that "in preattentive vision, precise positional relationships between textons are not important; only densities matter". The discrimination between textures will consist in averaging over the response of the dictionary elements.

In the example of the straw texture inserted into a paper texture, the discrimination is achieved using the information concerning the textons, and mainly on the energy density relative to the orientation and the scale. Indeed, the energy distribution of the paper is nearly constant at a given scale for different orientations, and is concentrated in the vertical "cells" at low and high scales and in the vertical and horizontal "cells" at intermediate scales.

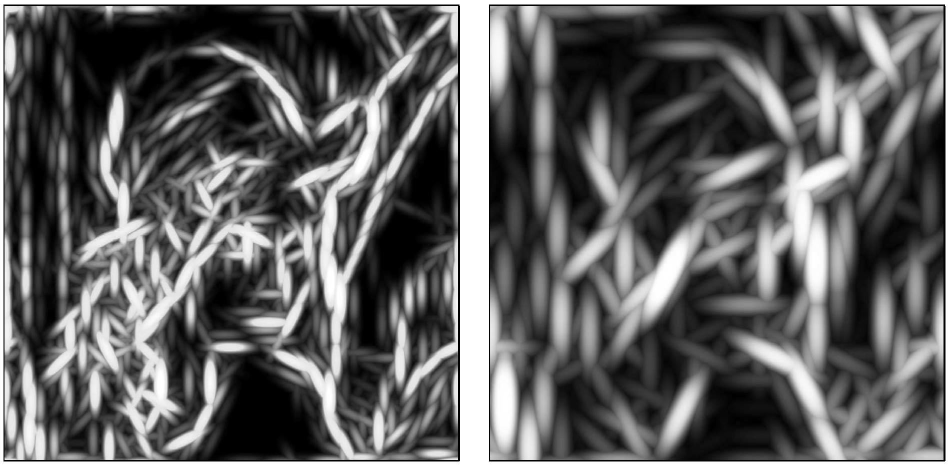

Figure 5. Symbolic representation of the Gabor vectors: the left and right figures corresponds respectively to the scales 2^3 and 2^4.

Figure 6. Reconstruction of the Lena image with 2500 (left figure) and 5000 atoms (right figure). The reconstruction is obtained as a truncated sum of the decomposition.

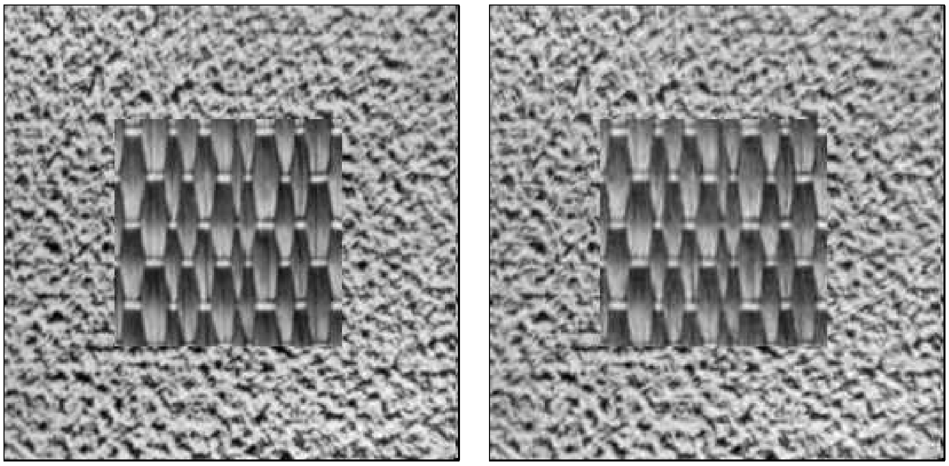

Figure 7. Original paper and straw texture mosaic (256×256, left image) and reconstructions with 5000 atoms (right image).

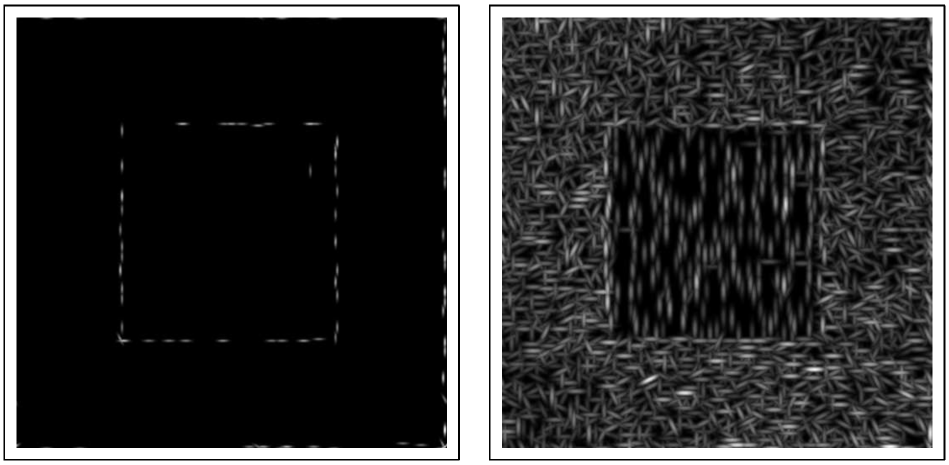

Figure 8. Symbolic representation of the Gabor vectors for the texture mosaic: scale 2^1 (left image) and 2^2 (right image).

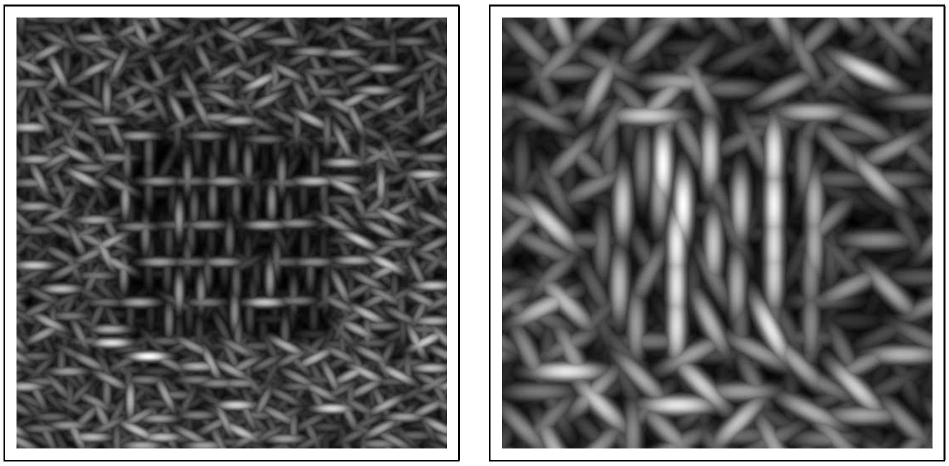

Figure 9. Symbolic representation of the Gabor vectors for the texture mosaic: scale 2^3 (left image) and 2^4 (right image).

§6 Vision applications

We introduced here a method to construct a decomposition of images into its main features. We showed that this transform provides a precise and complete characterization of the edges and texture components in terms of localization, orientation, scale and amplitude. By reconstructing high-visual quality images with very few atoms, we also showed that this representation is compact.

Another advantage of Matching Pursuit is the flexibility of the dictionary choice allowing to explicitely introduce *a priori* knowledge on the features of object classes into the dictionary to solve specific vision problems.

First experiments on image denoising suggest the use of a higher level information on edges obtained by atom linking to discriminate image features from noise.

A promising aspect of this method is its aptitude to model some early vision mechanisms in the visual cortex and especially those involved in texture discrimination. The Matching Pursuit with a Gabor dictionary acts as a linear filtering followed by a non-linearity (selection of the maximal coefficient) and some local inhibitions in the representation (update), which is very similar to the recent model of early vision mechanisms involved in texture discrimination and introduced by Malik and Perona ([11]).

As intrinsic steps of the Matching Pursuit, these two fundamental operations appears here in a very natural way as a consequence of the redundancy of the dictionary. First, the selection of the best matching element of the dictionary yields the amplitude of the vector, including its sign through the phase information. This corresponds to a full wave rectification, but contrary to some other work in the same field, is combined with the phase extraction and thus keeps the essential information (we still can discriminate the two symetrical textures, one composed of the repetition of a micro-pattern M, and the other of the repetition of the opposite micro-pattern -M). It is also biologically plausible since Pollen and Ronner ([13]) suggested the existence of cells pairs in quadrature phase in the cortex. Second, the update of the representation at each step locally removes the coefficients of the Gabor functions which do not correspond to an optimal match with the selected image feature. This mechanism appears to be very similar to the interactions among the neurons in the primary visual cortex, as described in ([11]). This model introduce the inhibition among the neurons as a simple consequence of the selection of the best responding neuron (the dictionary element in our model): once a predominant neuron is selected, all the spurious responses in its neighbourhood are removed, preventing the choice of neurons wich are not optimally tuned to the visual stimulus. This local enhancement of the best responding neuron is a variant of the winner-take-all mechanism previously introduced in the neural network litterature and hypothetized by Barlow: " Our perception are caused by the activity of a rather small number of neurons selected from a very large population of predominantly silent cells". Moreover, the intrinsic inhibition law of the matching pursuit gives an explicit formula for the interactions among neurons as the inner product of the dictionary elements, which does not depends on any parameter.

The decomposition of textured images proved the efficient discrimination properties of the method and suggests the computation of a texture gradient by an appropriate pooling of the gradient responses for each scale and orientation. However, instead of using a gradient computed with a smoothed version of the matching pursuit representation (cf. [11]), which assumes an arbitrarily fixed smoothness parameter, we are studying some clustering techniques based on the uniformity of the atom density, the predominant orientation and the amplitude. We then compute a gradient between these regions, without making any *a priori* assumption on their size.

Acknowledgements. This work was supported by the AFOSR grant F49620-93-1-0102, ONR grant N00014-91-J-1967 and the Alfred Sloan Foundation.

References

[1] J. Beck, "Textural segmentation", representation in perception, Beck J. ed., Organization and Lawrence Erlbaum Associates, Hillsdale, N.J., pp 285-317.

[2] J. Daugman, "Uncertainty relation for resolution in space, spatial-fequency and oriented optimized by two-dimensional visual cortex filters", *J. Opt. Soc. Am.*, vol. 2, pp 1160-1169, 1985.

[3] J. Daugman, "Complete discrete 2D Gabor transform by neural networks for image analysis and compression", *IEEE Trans. on Acoustic, Speech, and Signal Processing*, ASSP-36, pp 1169-1179, 1988.

[4] G. Davis, S. Mallat and M. Avenaleda, "Chaos in adaptive approximations", Technical Report, Computer Science, No. 325, April 1994.

[5] D. Gabor, "Theory of communication", *J. Inst. Elect. Ing.*, vol. 93, pp 429-457, 1946.

[6] P. J. Huber, "Projection Pursuit", *The Annals of Statistics*, vol. 13, No. 2, pp. 435-475, 1985.

[7] L. K. Jones, "On a conjecture of Huber concerning the convergence of projection pursuit regression", *The Annals of Statistics*, vol. 15, No. 2, pp. 880-882, 1987.

[8] J. Jones, L. Palmer, "An evaluation of two-dimensional Gabor filters model of simple receptive fields in cat striate cortex", *J. Neurophisiol.*, vol. 58, pp. 538-539, 1987.

[9] B. Julesz, J.R. Bergen, "Textons, the fundamental elements in preattentive vision and perception of textures", *The Bell System Tech. J.*, J. 62, pp. 1619-1645, 1983.

[10] S. Mallat and Z. Zhang, "Matching Pursuit with time-frequency dictionaries", *IEEE Trans. on Signal Processing*, Dec. 1993.

[11] J. Malik and P. Perona, "Preattentive texture discrimination with early vision mechanisms", *J. Opt. Soc. Am.*, vol. 7, No. 5, pp. 923-932, May 1990.

[12] S. Marcelja, "Mathematical description of the response of simple cortical cells", *J. Opt. Soc. Am.*, vol. 70, No. 11, pp 1297-1300, November 1980.

[13] D.A.Pollen, S.F. Ronner, "Visual cortical neurons as localized spatial filters", *IEEE Trans. on Systems, Man, and Cybernetics*, SMC-13, No. 5, pp 907-936, 1983.

[14] M.R. Turner, "Texture discrimination by Gabor functions", *Biol. Cybern.*, vol. 55, pp. 71-82, 1986,

[15] A.B. Watson, A.J. Ahumada, Jr., "A hexagonal orthogonal-oriented pyramid as a model of image representation in visual cortex", *IEEE Trans. Biomed. Eng.*, vol. 36, No.1, pp. 97-106, Jan. 1989.

François Bergeaud
Ecole Centrale Paris, Applied Mathematics Laboratory,
Grande Voie des Vignes, F-92290 Châtenay-Malabry, France
francois@mas.ecp.fr

Stéphane Mallat
Courant Institute of Mathematical Sciences, New York University
251, Mercer Street, New York, NY 10012
mallat@image.nyu.edu