



HAL
open science

Local Frequency Analysis with Two-Dimensional Wavelet Transform

Caroline Gonnet, Bruno Torr sani

► **To cite this version:**

Caroline Gonnet, Bruno Torr sani. Local Frequency Analysis with Two-Dimensional Wavelet Transform. Signal Processing, 1994, 37 (3), pp.389-404. 10.1016/0165-1684(94)90007-8 . hal-01739702

HAL Id: hal-01739702

<https://amu.hal.science/hal-01739702>

Submitted on 21 Mar 2018

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destin e au d p t et   la diffusion de documents scientifiques de niveau recherche, publi s ou non,  manant des  tablissements d'enseignement et de recherche fran ais ou  trangers, des laboratoires publics ou priv s.

LOCAL FREQUENCY ANALYSIS WITH TWO-DIMENSIONAL WAVELET TRANSFORM

Caroline Gonnet *

Bruno Torresani **

Abstract : An algorithm for the characterization of local frequencies in two-dimensional signals is described. The algorithm generalizes to the 2D situation a method previously proposed in the 1D context, based on the analysis of the phase of the continuous wavelet transform. It uses families of wavelets generated from a unique one by shifts, dilations and rotations.

Résumé: On décrit un algorithme permettant la caractérisation de fréquences locales dans les signaux bidimensionnels. L'algorithme est la généralisation bidimensionnelle d'une méthode précédemment proposée dans le contexte unidimensionnel, basée sur l'analyse de la phase de la transformée en ondelettes continue. Il utilise une famille d'ondelettes engendrées à partir d'une unique fonction par translations, dilatations et rotations.

Keywords: Wavelets, local frequencies, texture, asymptotic approximations.

22 pages, 14 figures.

January 1993; revised April 1993.

CPT-93/P.2870

* : ENSIEG Grenoble; supported by DIGILOG, Les Milles.

** : Groupement De Recherches Ondelettes, CNRS.

I. INTRODUCTION.

The problem of determining local frequencies in signals occurs in many different signal processing applications, in one-dimensional as well as two and three dimensional situations, for example in speech analysis and computer musics (1D) or machine vision (2D). It has been recognized for a long time that time-frequency methods (1D) or their higher dimensional generalizations can provide useful tools to handle such problems. Moreover, physiological data tend to show that image coding in our visual system could be realized by means of frequency and orientation selective cells.

Particularly interesting in this context are the continuous Gabor and wavelet decompositions, which provide a random access to the phase space (or position-frequency space), together with simple reconstruction formulas for retrieving the signal from its transform. Unfortunately, it is in general necessary to sample these transforms, so that position and frequency precision are somewhat lost. On the other hand, a very fine sampling of the continuous wavelet transform is difficult to use, because of prohibitive computational time (at least for the 2D and 3D cases). Nevertheless it has been recently shown that such sampled transforms can give useful informations about local frequencies in images (see for example [3] and [2]). In particular, the authors of [2] are able to extract local frequencies in images by considering the output of a small number of wavelet-like filters located in a neighborhood of the searched frequency. They show that a sufficiently dense sampling of the frequency plane with such filters yields a sufficient information to get an approximate value of the frequency.

Recently, a method was proposed in the one-dimensional case by one of us and coworkers [7] (see also [5] for a review of the theory and results in the one-dimensional case), that keeps the advantages of continuous decompositions and reduces the computational time. The main idea is that when the analyzed signal contains a component to which a local frequency can be associated, it is sufficient to study the wavelet transform to a neighborhood of a given curve (called the ridge of the wavelet transform) in the phase space to characterize the frequency. Moreover, there exists a fixed-point algorithm, based on some a-priori assumptions on the signal, yielding an estimate for this curve. It is important to point out that such algorithms are based on a careful study of the phase of the wavelet transform rather than the modulus (or phase-space energy localization), which seem to yield more precise estimates (see also [3]).

We describe here a two-dimensional generalization [8] of this one-dimensional algorithm. The notion of instantaneous frequency is replaced by the notion of local wavevector. To be able to characterize such local wavevectors with wavelets, it is necessary to use a version of wavelet analysis in which the frequency localization of wavelets can be controlled. This is done following [12] by adding an extra rotation degree of freedom, in addition to the usual scale and (two-dimensional) position parameters. One then gets a family of filters that are localized both in space and frequency. Their frequency localization is the same as that of the filters used in [2] and [3]. The structure of the algorithm yields a random access to the position-frequency space, and then a good precision in the extraction and characterization of frequency modulated components. The main difference with the existing methods is that our algorithm is adaptive, and not associated with an a-priori sampling of the frequency plane. Under suitable assumptions, it also provides

approximate expressions for the components of the signal associated with the extracted frequencies.

I.1. Two-dimensional wavelet transform.

Let us first describe the basics of the version of two-dimensional wavelet analysis we are interested in (see [12] for a detailed account of the method and some illustrations, and [1] for a detailed study of the performances of continuous wavelet decompositions in a wide range of different contexts).

Our convention for Fourier transform is the following one: $\hat{f}(\underline{\xi}) = \int f(\underline{x}) e^{i\underline{\xi} \cdot \underline{x}} d\underline{x}$, so that Plancherel's formula reads: $\|\hat{f}\| = 2\pi\|f\|$.

Let $\psi(\underline{x}) \in L^1(\mathbb{R}^2) \cap L^2(\mathbb{R}^2)$ be a *complex-valued* function, localized in such a way that:

- $\psi(\underline{x})$ is compactly supported in a neighborhood I_ψ of the origin, and is maximal at the origin $\underline{x} = 0$.
- $\hat{\psi}(\underline{\xi})$ is locally maximal at a given value of the frequency $\underline{\xi} = \underline{k}_0$, and has good decay properties far away from \underline{k}_0 . It can be assumed for instance that $\hat{\psi}(\underline{\xi}) \leq K \cdot \|\underline{\xi} - \underline{k}_0\|^{-\alpha}$ for some constant K and some $\alpha > 2$ (notice that this implies that $\psi(\underline{x})$ is at least of class $C^3(\mathbb{R}^2)$).
- The following constant is finite and nonzero

$$c_\psi = \int_{\mathbb{R}^2} |\hat{\psi}(\underline{\xi})|^2 \frac{d\underline{\xi}}{\|\underline{\xi}\|^2} \quad (I.1)$$

Notice that the non-vanishing of c_ψ implies that $\hat{\psi}(0) = \int_{\mathbb{R}^2} \psi(\underline{x}) d\underline{x} = 0$. Since $\psi \in L^2(\mathbb{R}^2)$, $\hat{\psi}(\underline{\xi}) \rightarrow 0$ when $\underline{\xi} \rightarrow \infty$, and ψ may then be thought of as a band-pass filter.

Such a function will be called an *analyzing wavelet*. Associate with it the following family of wavelets, shifted, scaled and rotated copies of $\psi(\underline{x})$:

$$\psi_{(\underline{b}, a, \theta)}(\underline{x}) = \frac{1}{a^2} \psi\left(\underline{r}_{\theta}^{-1} \cdot \frac{\underline{x} - \underline{b}}{a}\right) \quad (I.2)$$

where

- $\underline{b} \in \mathbb{R}^2$ is a translation parameter,
- $a \in \mathbb{R}_+^*$ is a scale (or dilation) parameter,
- $\theta \in [0, 2\pi[$ is a rotation parameter, and
- \underline{r}_{θ} is the rotation matrix of angle θ :

$$\underline{r}_{\theta} = \begin{pmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{pmatrix} \quad (I.3)$$

Any $f(\underline{x}) \in L^2(\mathbb{R}^2)$ can be decomposed as a superposition of such elementary wavelets, as

$$f(\underline{x}) = \frac{1}{c_\psi} \int_{\mathbb{R}^2 \times \mathbb{R}_+^* \times [0, 2\pi[} T_f(\underline{b}, a, \theta) \psi_{(\underline{b}, a, \theta)}(\underline{x}) d\underline{b} \frac{da}{a} d\theta \quad (I.4)$$

Here, the coefficients $T_f(\underline{b}, a, \theta)$ appearing in the decomposition are nothing but the scalar products of the analyzed function $f(\underline{x})$ with the wavelets $\psi_{(\underline{b}, a, \theta)}(\underline{x})$:

$$T_f(\underline{b}, a, \theta) = \int f(\underline{x}) \psi_{(\underline{b}, a, \theta)}(\underline{x})^* d\underline{x} = \frac{1}{a^2} \int_{\mathbb{R}^2} f(\underline{x}) \psi\left(\frac{\underline{x} - \underline{b}}{a}\right)^* d\underline{x} \quad (I.5)$$

The family of wavelets can then be used almost as if it formed an orthonormal basis of $L^2(\mathbb{R}^2)$, which is far from being the case, since this family is overcomplete. This means that the wavelets are linearly independent. The algorithms described below are based on this redundancy.)

Remark: The assumption that the wavelet $\psi(\underline{x})$ is compactly supported is useful as far as mathematical estimates are concerned, but not necessary from the numerical point of view. Indeed the algorithms described here give results with an accuracy of the order of about 3 digits. Then if $\psi(\underline{x})$ is not compactly supported, it is enough to assume that it has fast decay (i.e. exponential decay, or algebraic decay with large exponent), and the numerical values of the wavelet smaller than the given accuracy are neglected. The same remark also applies to the admissibility of the wavelets. Since one works with a fixed accuracy, it is sufficient to assume that $\hat{\psi}(0)$ is smaller than the desired accuracy; such wavelets are sometimes called "numerically admissible".

I.2. Energy localization in phase space.

Let $\psi(\underline{x})$ be an analyzing wavelet as described in the previous section, and let $f(\underline{x})$ be a signal to be analyzed. Then due to the localization properties of the wavelet $\psi(\underline{x})$, the coefficient $T_f(\underline{b}, a, \theta)$ gives informations on the "content of $f(\underline{x})$ " near the point $\underline{x} = \underline{b}$. Moreover, since Plancherel's formula yields

$$T_f(\underline{b}, a, \theta) = \frac{1}{2\pi} \langle \hat{f}, \widehat{\psi_{(\underline{b}, a, \theta)}} \rangle = \frac{1}{2\pi} \int_{\mathbb{R}^2} \hat{f}(\underline{\xi}) e^{i\underline{\xi} \cdot \underline{b}} \hat{\psi}(a\underline{r}_{-\theta} \cdot \underline{\xi})^* d\underline{\xi} \quad (I.6)$$

so that the wavelet coefficient $T_f(\underline{b}, a, \theta)$ also measures the "content of $\hat{f}(\underline{\xi})$ " near the frequency $\underline{\xi} = \underline{r}_{-\theta} \cdot \underline{k}_0/a$.

In other words, the wavelet $\psi_{(\underline{b}, a, \theta)}$ is localized in the neighborhood of the point $(\underline{b}; \underline{r}_{-\theta} \cdot \underline{k}_0/a)$ in the *phase-space* (or *position-spatial frequency space*), and the corresponding wavelet coefficient $T_f(\underline{b}, a, \theta)$ "gives a measure" of the "phase-space content" of $f(\underline{x})$ near that point. Notice that (a^{-1}, θ) appear as the polar coordinates of a frequency naturally associated with $\psi_{(\underline{b}, a, \theta)}(\underline{x})$. The frequency localization of the wavelets is described in figure 1.

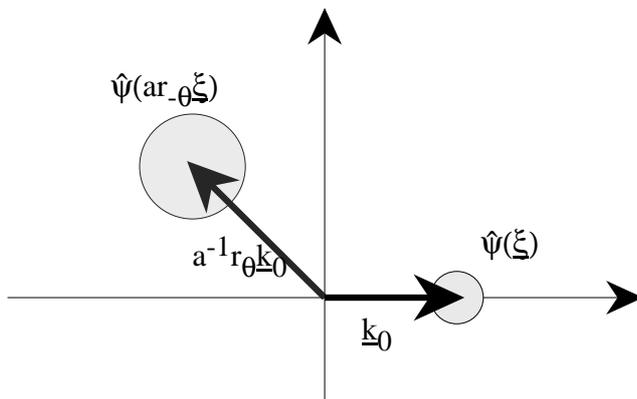


Figure 1: Localization of wavelets in the Fourier space

Warning: It is crucial to notice that there is no completely natural notion of "phase-space content" of a function, as a consequence of Heisenberg's uncertainty principle. This is illustrated by the fact that the wavelet coefficient of course depends on the choice of the chosen analyzing wavelet $\psi(\underline{x})$. The interpretation of the wavelet transform has then to take into account the properties of the chosen wavelet. In practice, it is important to choose an analyzing wavelet sufficiently well-localized in the position-frequency space. It is well-known that gaussian functions minimize the Heisenberg inequality (see also [2] for a careful analysis of this problem in a context similar to that of the present paper). It is then particularly interesting to use analyzing wavelets with a Gaussian amplitude. In the numerical examples that we consider, we shall always use analyzing wavelets of the form $\exp\{-\underline{x}^2\} \exp\{i\underline{k}_0 \cdot \underline{x}\}$ (Morlet wavelets), and $\exp\{-\underline{x}^2\} \exp\{i(k_0 x + \alpha x^2/2)\}$ (chirp wavelets). The reason why we need the second choice will be explained later.

A simple illustration of the phase-space localization properties of wavelet analysis is provided by the following simple example. Consider a function of the form

$$f(\underline{x}) = A(\underline{x})e^{i\underline{k} \cdot \underline{x}} \quad (I.7)$$

where $A(\underline{x})$ is some (slowly-varying) amplitude function, assumed to be at least once differentiable. Then using Taylor's formula for expanding $A(\underline{x})$ around $\underline{x} = \underline{b}$, we get the following approximate expression for the wavelet coefficient

$$T_f(\underline{b}, a, \theta) = A(\underline{b})e^{i\underline{k} \cdot \underline{b}} \hat{\psi}(a \underline{r}_{-\theta} \cdot \underline{k})^* + R(\underline{b}, a, \theta) \quad (I.8)$$

Here $R(\underline{b}, a, \theta)$ is a remainder, that can be estimated as follows (see e.g. [9]):

$$|R(\underline{b}, a, \theta)| \leq a^2 \sup_{a^{-1}\underline{x} \in I_\psi} \|\nabla A(\underline{x})\| \quad (I.9)$$

Here ∇ stands for the gradient operator.

The presence of the term $\hat{\psi}(a\underline{r}_{-\theta} \cdot \underline{k})^*$, together with the localization properties of $\hat{\psi}(\underline{\xi})$ show that the wavelet coefficient $T_f(\underline{b}, a, \theta)$ is small, except in a neighborhood of the point $(\underline{b}; a^{-1}\underline{r}_{-\theta} \cdot \underline{k}_0)$ in the phase space, where the wavelet transform is locally maximal.

If the signal to be analyzed is now a superposition of such components, of the form

$$s(\underline{x}) = \sum_{i=1}^N A_i(\underline{x}) e^{i\underline{k}_i \cdot \underline{x}} , \quad (I.10)$$

the linearity of the wavelet transform implies that $T_f(\underline{b}, a, \theta)$ is the sum of the wavelet transforms of all the components, i.e. is (up to some remainder) the sum of terms

$$A_i(\underline{b}) e^{i\underline{k}_i \cdot \underline{b}} \hat{\psi}(a\underline{r}_{-\theta} \cdot \underline{k}_i)^* , \quad (I.11)$$

that have in general different phase-space localization properties (i.e. are localized in different regions of the phase space). Then if for some value of $(\underline{b}, a, \theta)$ one has $A_i(\underline{b}) \hat{\psi}(a\underline{r}_{-\theta} \cdot \underline{k}_i)^* \gg A_j(\underline{b}) \hat{\psi}(a\underline{r}_{-\theta} \cdot \underline{k}_j)^* \forall j \neq i$ then $T_f(\underline{b}, a, \theta) \cong A_i(\underline{b}) \hat{\psi}(a\underline{r}_{-\theta} \cdot \underline{k}_i)^*$, which means that the wavelet transform has (up to some accuracy) selected one component.

II. MODEL AND APPROXIMATIONS.

As stressed in the introduction, the approach we develop is based on a local version of spectral analysis, realized via the two-dimensional wavelet transform.

II.1. The model.

We model a regular texture pattern as a superposition of components of the form

$$s_i(\underline{x}) = A_i(\underline{x}) \cos(\phi_i(\underline{x})) \quad (II.1)$$

where $A_i(\underline{x})$ is a real-valued amplitude function, assumed to vary slowly compared with the oscillatory term $\cos(\phi_i(\underline{x}))$, and the index i runs over the different components of the signal $s(\underline{x})$. We will specify a little bit more the necessary assumptions later on.

Such a model is very classical in one-dimensional signal processing and in physical applications (in optics for instance). It was also used as image model in [2] (and maybe implicitly in [3]). It is worth noticing that (as stressed in [2]) such a model is far from unique, and only represents a possible description of images. Another point to be stressed is that images in general only take positive values, in contrast with the model (II.1). Nevertheless, since wavelets are of zero integral, they operate modulo constants, so that the results of the analysis does not change if we add to (II.1) some positive constant in order to make the model more realistic for images.

Let us now consider an analyzing wavelet $\psi(\underline{x})$, such that its Fourier transform $\hat{\psi}(\underline{\xi})$ is maximal at the frequency $\underline{\xi} = \underline{k}_0$. Then the wavelet transform $T_s(\underline{b}, a, \theta)$ is locally maximum around the phase-space points (see below)

$$(\underline{b}; a_i = \frac{\|\underline{k}_0\|}{\|\nabla\phi_i(\underline{b})\|}, \theta_i = \frac{\langle \underline{k}_0, \nabla\phi_i(\underline{b}) \rangle}{\|\underline{k}_0\| \|\nabla\phi_i(\underline{b})\|}) \quad (II.2)$$

The main idea of the algorithm we are going to describe is to restrict the analysis of the wavelet transform to a neighborhood of such phase-space points, in order to characterize the corresponding component. The same procedure will then be applied to all components.

The first step is then to focus to a given component, and to characterize it from its wavelet transform. To do so, we shall use some approximate expressions for the wavelet transform of signals we are interested in.

II.2. Some approximate formulas.

We have seen in section I.2 an example of approximate formula for the wavelet transform of an academic signal. Such a formula can be generalized to the academic signals discussed in section II.1, under the additional assumption that the phase function $\phi(\underline{x})$ is slowly varying. For instance, it may easily be shown (still using Taylor developments, see e.g. [5] in the one-dimensional case, and in a slightly different form [2] in the image analysis context) that the wavelet transform of $A \exp\{i\phi\}$ may be approximated as

$$T_s(\underline{b}, a, \theta) = A(\underline{b})e^{i\phi(\underline{b})}\hat{\psi}(a\underline{r}_{-\theta} \cdot \nabla\phi(\underline{b}))^* + R(\underline{b}, a, \theta)$$

However, more precise approximations can be obtained by means of the stationary phase method (or related approximations). For an overview of the method, we refer to [6], [9].

To describe the method, let us start with an academic signal

$$s(\underline{x}) = A_i(\underline{x}) \cos(\phi(\underline{x})) \quad (II.3)$$

and consider an analyzing wavelet, put on the form

$$\psi(\underline{x}) = A_\psi(\underline{x})e^{i\phi_\psi(\underline{x})} \quad (II.4)$$

The wavelet coefficient $T_s(\underline{b}, a, \theta)$ then takes the form of an oscillatory integral:

$$\begin{cases} T_s = \frac{1}{a^2} \int A(\underline{x})A_\psi\left(\underline{r}_{-\theta} \cdot \frac{\underline{x}-\underline{b}}{a}\right)e^{i\Phi(\underline{b}, a, \theta)(\underline{x})} d\underline{x} \\ \Phi(\underline{b}, a, \theta)(\underline{x}) = \phi(\underline{x}) - \phi_\psi\left(\underline{r}_{-\theta} \cdot \frac{\underline{x}-\underline{b}}{a}\right) \end{cases} \quad (II.5)$$

Following the stationary phase argument, the main contribution to such an integral comes from the stationary points \underline{x}_s of the phase $\Phi(\underline{b}, a, \theta)(\underline{x})$, i.e. the points $\underline{x}_s = \underline{x}_s(\underline{b}, a, \theta)$ such that

$$\nabla\phi(\underline{x}_s) = \frac{1}{a}\underline{r}_{-\theta} \cdot \phi_\psi\left(\underline{r}_{-\theta} \cdot \frac{\underline{x}-\underline{b}}{a}\right)(\underline{x}_s) \quad (II.16)$$

A second order development of the phase $\Phi(\underline{x}) = \Phi(\underline{b}, a, \theta)(\underline{x})$ yields the following approximate expression (see appendix A):

$$\begin{cases} T_s(\underline{b}, a, \theta) = \frac{s(\underline{x}_s)\psi\left(\underline{r}_{-\theta} \cdot \frac{\underline{x}_s-\underline{b}}{a}\right)}{Corr(\underline{b}, a, \theta)} + R(\underline{b}, a, \theta) \\ Corr(\underline{b}, a, \theta) = \sqrt{|\partial_1^2\Phi \partial_2^2\Phi - (\partial_1\partial_2\Phi)^2|} e^{-i\frac{\pi}{2} \text{sgn}(\partial_1^2\Phi \partial_2^2\Phi - (\partial_1\partial_2\Phi)^2)} \end{cases} \quad (II.7)$$

where $R(\underline{b}, a, \theta)$ is a remainder, that can be estimated using standard techniques [9] ($R(\underline{b}, a, \theta)$ actually depends on the derivatives of the amplitudes up to order 2, and the derivatives of the phases up to order 4). Notice that such an approximation is slightly more precise than the previous one, but the form of the remainder is much more involved and difficult to handle. Also, the first approximation yields an exact solution in the case of an harmonic signal analyzed with a Morlet wavelet, whereas the stationary phase one is singular. In such cases, the wavelet can be modified so that the singularity is avoided, for instance by choosing chirp wavelets instead of Morlet wavelets.

Remark: To derive such an expression, we had to assume that for the considered values of $(\underline{b}, a, \theta)$, there exists a unique stationary point $\underline{x}_s(\underline{b}, a, \theta)$. We shall briefly discuss later what happens when this is not the case.

II.3. Ridge and skeleton of the wavelet transform:

As in the one-dimensional case, one clearly sees on such an expression that the information concerning $s(\underline{x})$ is completely contained in a subset of the phase space, called the *ridge of the wavelet transform*. Indeed, due to the localization of the analyzing wavelet in the \underline{x} -space, $T_s(\underline{b}, a, \theta)$ is essentially localized around the phase-space points $(\underline{b}, a, \theta)$ such that $\underline{x}_s(\underline{b}, a, \theta) = \underline{b}$, if such points do exist. We then introduce the ridge:

$$\mathcal{R} = \{(\underline{b}, a, \theta) \in \Omega \text{ s.t. } \underline{x}_s(\underline{b}, a, \theta) = \underline{b}\} \quad (II.8)$$

Plugging this definition into that of the stationary points yields

$$\nabla\phi(\underline{b}) = a^{-1} \underline{r}_{\underline{\theta}} \cdot \nabla\phi_\psi(0) \quad \text{on the ridge.} \quad (II.9)$$

or, otherwise stated:

$$\nabla\phi(\underline{x}) = a_r(\underline{x})^{-1} \underline{r}_{\underline{\theta}_r(\underline{x})} \cdot \nabla\phi_\psi(0) \quad (II.10)$$

Remark: The ridge of the wavelet transform then appears as a two-dimensional surface in the four-dimensional phase space. More precisely, it takes the form of a vector field $\underline{k}_r(\underline{b})$, of polar coordinates $(a^{-1} = a_r(\underline{b})^{-1}, \theta = \theta_r(\underline{b}))$. Again the inverse-scale and rotation parameters can be interpreted as polar coordinates of local wavevectors.

To show that the signal $s(\underline{x})$ is completely determined by its restriction to the ridge (as in the one-dimensional case, we shall call this restriction the *skeleton of the wavelet transform*), we have to check that the restriction of the correction function $Corr(\underline{b}, a, \theta)$ to the ridge is completely determined by the ridge itself. This can be directly seen from the expression of the correction function (see appendix B). Then the expression of the skeleton of the wavelet transform reads:

$$T_s(\underline{b}, a_r(\underline{b}), \theta_r(\underline{b})) = \frac{2\pi \psi(0)^*}{Corr(\underline{b}, a_r(\underline{b}), \theta_r(\underline{b}))} s(\underline{b}) \quad (II.11)$$

which implies that the knowledge of the ridge and the skeleton of the wavelet transform is sufficient to characterize $s(\underline{x})$.

Remark: As shown in Appendix A.1, the error estimate provided by the stationary phase method involves in particular the derivatives of the amplitudes $A(\underline{x})$ and $A_\psi(\underline{x})$, or more precisely their logarithmic derivatives. In other words, the slowest the variations of the amplitudes compared with the oscillations, the best the approximation. But it is clearly better (at least for practical purpose) to use wavelets that do not have such a property, i.e. analyzing wavelets that do not oscillate too much. It turns out that in the case of analyzing wavelets with a Gaussian amplitude function (i.e. with $A(\underline{x}) = \exp\{-\underline{x} \cdot \underline{A} \cdot \underline{x}\}$ where \underline{A} is a positive-definite symmetric matrix), the previous approximation can be slightly modified in order to take into account the exact form of the amplitude. The resulting approximate expression is derived in appendix A.2, and has roughly the same form as the previous one (the only change is in the correction function). It is then still possible to introduce a ridge and a skeleton. The algorithms described in the next section can then be adapted to such a situation.

II.4. Multicomponent signals:

The case of multicomponent signals is in general a difficult problem, since different components can have different behaviours (they can for instance merge, or die suddenly, ...). Classical methods fail in such cases, and our algorithm will fail as well (because the approximations it is based on will no longer be valid). Nevertheless, since wavelet analysis is a local method, the approximations remain valid outside some neighborhood of such singular regions. Then the multicomponent problem can be handled in the same way as before, under the assumption that "the wavelet separates the ridges", i.e. that the ridges are far away enough from each other so that the contribution of the other components to a given skeleton can be neglected. This is quite easy to understand in the simpler approximation: for a signal modelet as in (II.1),

$$T_s(\underline{b}, a, \theta) \cong \sum_{k=1}^N A_k(\underline{b}) e^{i\phi_k(\underline{b})} \hat{\psi}(a \underline{r}_{-\theta} \cdot \nabla \phi_k(\underline{b}))^* \quad (II.12)$$

If at some point $(\underline{b}, a, \theta)$ the wavelet $\psi(\underline{x})$ is sufficiently well localized in frequency so that

$$A_k(\underline{b}) |\hat{\psi}(a \underline{r}_{-\theta} \cdot \nabla \phi_k(\underline{b}))| \gg A_l(\underline{b}) |\hat{\psi}(a \underline{r}_{-\theta} \cdot \nabla \phi_l(\underline{b}))| \quad \forall l \neq k$$

the wavelet transform will only "see" the contribution of $A_k(\underline{b}) e^{i\phi_k(\underline{b})}$, and will be able to characterize it. In such a case, if the analyzed signal is a superposition of N components of the type $A_i e^{i\phi_i}$, the wavelet transform exhibits N separate ridges, from which the N components can be extracted and reconstructed separately.

The cases where there are interacting ridges are in general much more difficult to handle, and the use of the algorithm described here leads to biased estimates for the local frequencies and amplitudes. Such cases have been studied in details in the one-dimensional case on [4] (see also [5]), but without providing a method for completely solving the problem.

The same kind of remark apply in the case of boundaries. Since in such cases our approximations are far from valid, the estimates that we obtain in the neighborhoods of

boundaries are biased (this is visible in the boundaries of the figures in the section devoted to numerical results). To precisely characterize such boundaries, it is necessary to use other methods (wavelet-based methods have been developed in different contexts, see e.g. [11]).

III. ALGORITHMS AND IMPLEMENTATION.

As we saw in the last section, all the needed information can be extracted from the ridge (or the ridges) of the wavelet transform. Then it is not necessary (for the kind of 2D signals we are interested in of course) to compute the whole wavelet transform, which would be very time-consuming. What is needed now is an algorithm that is able to give a reasonable estimate of the ridge (i.e. the local wavevector field) with a limited number of operations.

The algorithm we describe is just a two-dimensional version of the one-dimensional ridge extraction algorithms described in [5]. A simple possibility would be to use the localization properties of the wavelet in both the \underline{x} -space and the Fourier space; indeed, it follows from the approximate expression derived in the previous section that the modulus of the wavelet transform is locally maximum in the neighborhood of the associated ridge. Unfortunately the existence of the corrective term $Corr(\underline{b}, a, \theta)$ implies that this leads to a biased estimate of the ridge. Generalizing the algorithm proposed in the one-dimensional case, we found more precise to focus on the study of the phase of the wavelet transform.

III.1. Behaviour of the phase of the wavelet transform.

As a consequence of the previous estimates, the phase of the wavelet transform of $A(\underline{x}) \exp\{i\phi(\underline{x})\}$ is given by

$$Arg[T_s(\underline{b}, a, \theta)] = \phi(\underline{x}_s) - \phi_\psi\left(\underline{r}_{-\theta} \cdot \frac{\underline{x}_s - \underline{b}}{a}\right) - \frac{\pi}{2} \operatorname{sgn}(\partial_1^2 \Phi \partial_2^2 \Phi - (\partial_{12}^2 \Phi)^2) \quad (III.1)$$

A direct computation shows that

$$\nabla_{\underline{b}} Arg[T_s(\underline{b}, a, \theta)] = a^{-1} \underline{r}_{-\theta} \cdot \nabla \phi_\psi\left(\underline{r}_{-\theta} \cdot \frac{\underline{x}_s - \underline{b}}{a}\right) \quad (III.2)$$

hence, in particular

$$\nabla_{\underline{b}} Arg[T_s(\underline{b}, a, \theta)] = a^{-1} \underline{r}_{-\theta} \cdot \nabla \phi_\psi(0) \quad \text{on the ridge} \quad (III.3)$$

Then on the ridge, the local wavevector associated with the wavelet transform (i.e. the gradient of the phase) coincides with that of the rotated and scaled wavelet. We will use this fact as a characterization of the ridge.

Remark: In the case of a fixed-frequency wavelet such as the Morlet wavelet for example, $\nabla \phi_\psi$ is a constant, so that $\nabla \phi_\psi\left(\underline{r}_{-\theta} \cdot \frac{\underline{x}_s - \underline{b}}{a}\right) = \nabla \phi_\psi(0)$ everywhere. Then the usual

stationary phase approximation yields $\nabla_{\underline{b}} \text{Arg}[T_s(\underline{b}, a, \theta)] = a^{-1} \underline{r}_{-\theta} \cdot \nabla \phi_\psi(0)$ everywhere. In such a case it is better to use a more precise approximation (see appendix A.2), for which

$$\nabla_{\underline{b}} \text{Arg}[T_s(\underline{b}, a, \theta)] = a^{-1} \underline{r}_{-\theta} \cdot \nabla \phi_\psi\left(\underline{r}_{-\theta} \cdot \frac{\underline{x}_s - \underline{b}}{a}\right) + O\left(\underline{r}_{-\theta} \cdot \frac{\underline{x}_s - \underline{b}}{a}\right) \quad (III.4)$$

the additional term being in general nonzero. Then even in such a case, the above criterion can still be used as a characterization of the ridge.

III.2. A fixed-point algorithm.

Let us turn now to an algorithm based on that criterion. Owing to the above discussion, the problem is to find the phase-space points where $\nabla_{\underline{b}} \text{Arg}[T_s(\underline{b}, a, \theta)] = a^{-1} \underline{r}_{-\theta} \cdot \nabla \phi_\psi(0)$. We will make use of the following heuristic consideration: roughly speaking, the gradient of the phase of the wavelet transform results from a competition between the phase of the signal and the phase of the corresponding shifted, scaled and rotated wavelets. One may then expect it to lie in between the local wavevector of the signal and the gradient of the phase of the shifted, scaled and rotated wavelet. This suggests the following iterative scheme to reach it:

For a fixed value of the position parameter $\underline{b} = \underline{b}_0$, let $(a = a_0, \theta = \theta_0)$ correspond to an initial guess for a local frequency of the signal. Then the gradient of the phase of the wavelet transform $T_s(\underline{b}; a_0, \theta_0)$ gives another value for a local wavevector, whose polar coordinates are given by:

$$\begin{cases} a_1 = \frac{|\nabla \phi_\psi(0)|}{|\nabla \text{Arg}[T_s](\underline{b}_0, a_0, \theta_0)|} \\ \theta_1 = \arctan \frac{\langle \nabla \text{Arg}[T_s](\underline{b}_0, a_0, \theta_0), \nabla \phi_\psi(0)^\perp \rangle}{\langle \nabla \text{Arg}[T_s](\underline{b}_0, a_0, \theta_0), \nabla \phi_\psi(0) \rangle} \end{cases} \quad (III.5)$$

where $\nabla \phi_\psi(0)^\perp$ is defined as follows: $\nabla \phi_\psi(0)^\perp \perp \nabla \phi_\psi(0)$ and $\|\nabla \phi_\psi(0)^\perp\| = \|\nabla \phi_\psi(0)\|$.

Then iterate the procedure, until the obtained precision is smaller than a given (fixed) accuracy.

$$\begin{cases} a_{n+1} = \frac{|\nabla \phi_\psi(0)|}{|\nabla \text{Arg}[T_s](\underline{b}_0, a_n, \theta_n)|} \\ \theta_{n+1} = \arctan \frac{\langle \nabla \text{Arg}[T_s](\underline{b}_0, a_n, \theta_n), \nabla \phi_\psi(0)^\perp \rangle}{\langle \nabla \text{Arg}[T_s](\underline{b}_0, a_n, \theta_n), \nabla \phi_\psi(0) \rangle} \end{cases} \quad (III.6)$$

Finally, for the following values of \underline{b} , it is sufficient to take as initial guess $(a_0(\underline{b}), \theta_0(\underline{b}))$ the values of convergence for the previous position. The algorithm then continuously follows the ridge (within a given accuracy).

Remark: Since the problem is here a two-dimensional problem, it is necessary to specify a given direction (the x or the y direction in the case of sampled signals). The algorithm will then scan the x direction first, for fixed y parameter, and then increment the y parameter.

III.3. Reconstruction from the skeleton.

Once the ridge of the wavelet transform has been characterized by means of the previous fixed-point algorithm, the corresponding component can be reconstructed using the approximate expression of the skeleton of the wavelet transform:

$$s_{rec}(\underline{x}) = \frac{Corr(\underline{x}, a_r(\underline{x}), \theta_r(\underline{x}))}{2\pi \psi(0)^*} T_s(\underline{x}, a_r(\underline{x}), \theta_r(\underline{x})) \quad (III.7)$$

As stressed previously, the correction function $Corr(\underline{x}, a_r(\underline{x}), \theta_r(\underline{x}))$ is completely characterized by the ridge. Then for the reconstruction of the signal $s(\underline{x})$, the information provided by the ridge and the skeleton $T_s(\underline{x}, a_r(\underline{x}), \theta_r(\underline{x}))$ is sufficient .

III.4. Some remarks.

- The proposed algorithm is not the only one capable of giving an estimate for the ridge. It is also possible to use other procedures, such as for instance a Newton method for searching the zeroes of $\nabla Arg[T_s] - a^{-1} \underline{r}_0 \cdot \nabla \phi_\psi(0)$ [10]. The computational time seems a bit smaller in such a case.
- The algorithm avoids the computation of the whole wavelet transform (a function of four variables !) and then is a priori faster than expected. Nevertheless, in the case of large 2D signals, the computational time is still a problem. The problem can be avoided as follows. The method is based on approximations, with in particular the a-priori assumptions that the amplitude and the local wavevectors are slowly varying (this is necessary in order to neglect the remainders in the approximations). Then it makes sense to subsample the local amplitudes and wavevectors. Since the number of operations needed to process an $N \times N$ signal is proportional to N^2 , a subsampling rate of τ yields a division of the computational time by a factor of τ^2 . The fixed-point algorithm can then be run with a coarser sampling for the \underline{b} variable, yielding a coarser sampling of the local amplitudes and wavevectors, for which the finer samplings can be retrieved using an interpolation procedure. The subsampling rate τ is of course to be matched to the signal, more precisely to the speed of variation of the local amplitude and frequency. The reconstruction can then be performed from the fine sampling of the skeleton.

IV. NUMERICAL RESULTS.

To illustrate the performances of the method, let us consider some simple numerical examples. All the computations shown here have been done with the so-called chirp wavelet, with Gaussian amplitude (and phase term) in order to optimize the position-frequency localization:

$$\psi(x, y) = e^{-(x^2+y^2)} e^{ik_0(x+\alpha x^2/2)}$$

which is localized near the frequency $\underline{k}_0 = k_0 \underline{e}_1$ in the Fourier space. The initial guess for the a and θ parameters were determined visually from the image (there is no need of great precision for such initial guesses). The required precision was set to 10^{-3} , but a cutoff for

the maximal number of iterations (5 or 10 in general) was introduced in order to reduce computational time if necessary.

The first example (128×128 image) is provided by a superposition of two "chirped" components (figure 2.a), i.e. such that their local wavevectors depend linearly on the coordinates, amplitude-modulated by a Gaussian function. The wavelet-based algorithm is clearly able to separate the components, as shown in figure 2.b.

The local wavevectors (or wave directors, since they have no sign, the image being real-valued) are represented by scaled and rotated needles. The wavevectors have been computed only in regions of the image where the corresponding component is numerically significant. The reconstructions of the two corresponding components are shown in figure 3. The quality of the reconstruction is very good (there is no boundary effect because of the gaussian amplitude of the two components).

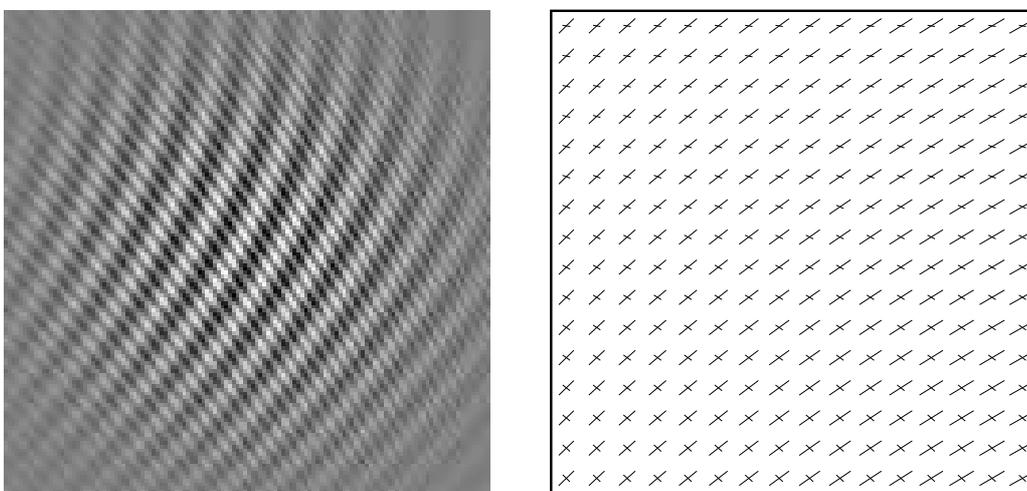


figure 2: sum of two chirped components: source image (a) and ridge (b)

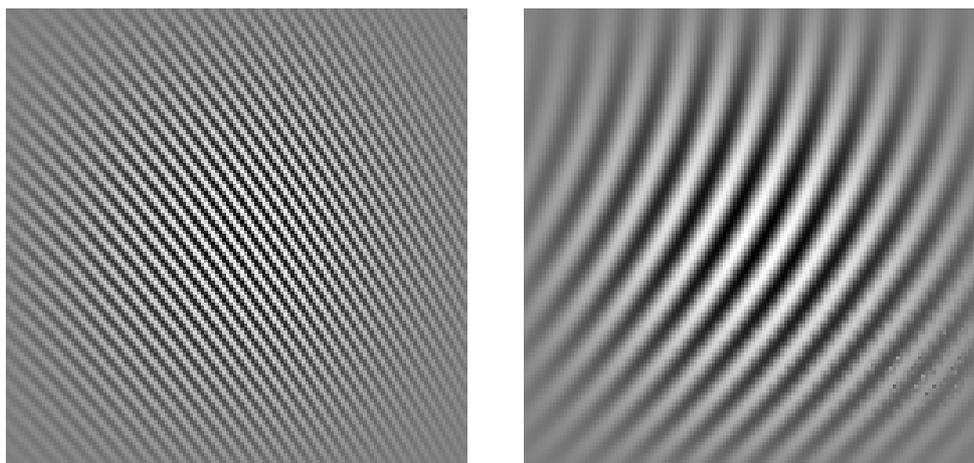


figure 3: Reconstruction of the two components from the skeleton

We also consider (see figure 4) another academic 128×128 texture image (referred to as the "elliptic example"), in which the local wavevectors' amplitude depend linearly on the coordinates, and their orientation depend linearly on the angle to a fixed axis (with respect to a origin fixed at the center of the image).

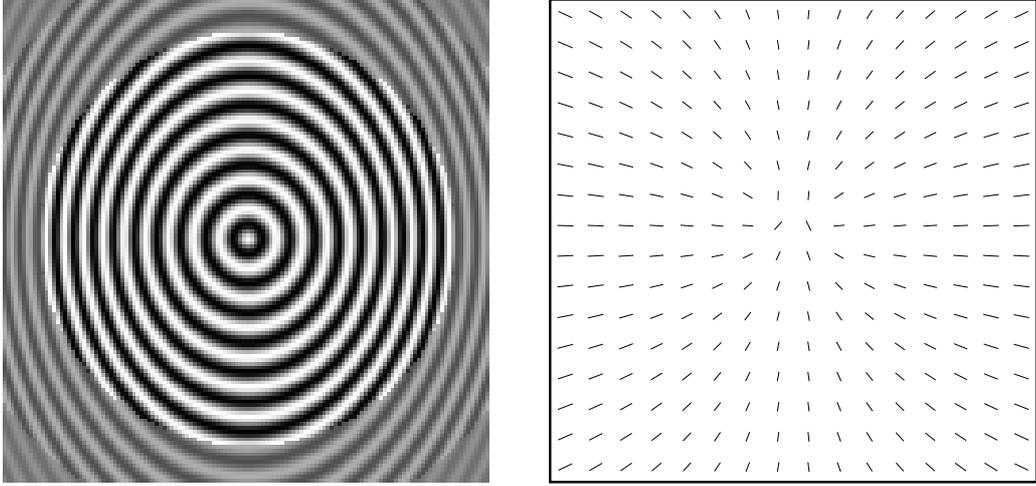


figure 4: "elliptic example": source image (a) and ridge (b)

Once again the characterization of the local wavevectors and the amplitude of the reconstruction are of quite good quality (figure 5), whereas the phase suffers from the existence of a topological defect at the center of the image (the local wavevector is not defined there). The discontinuities that appear in figure 5 in the phase and reconstruction images are easy to explain. The initial value for the fixed-point algorithm at (b_1, b_2) is taken to be the value of convergence at $(b_1, b_2 - \delta b)$ (except at the top and bottom boundaries). On a vertical line that is always far away from the center of the image (with respect to the size of the wavelet), the local wavevector can rotate around the center. In contrast, for a vertical line that comes near the center, the local wavevector should switch from \underline{k} to $-\underline{k}$ when crossing the center, which is not allowed by the algorithm. Instead, the algorithm switches itself to the negative frequency part of the (real-valued) signal.

Another simple example is provided by the case of regular texture shown in figure 6, i.e. a collection of T and L symbols, with various orientations. In such a case, the algorithm looks for some periodicity, and effectively extracts local wavevectors corresponding to periodicities present in the source image. Notice that it is difficult here to separate the different regions of the image if only the local wavevector information is kept (except the upper right corner, and the lower part of the image). Nevertheless, it can be seen that the phase and modulus of the wavelet transform (at various frequencies, or at a fixed frequency) can be used to complete the information (see [3] for example). For example, it is well known (see [11], [1] for example) that the modulus of the wavelet transform (for small values of the scale parameter) yields informations on the singularities of the image (in this cases, the singularities are the boundaries of the regions).

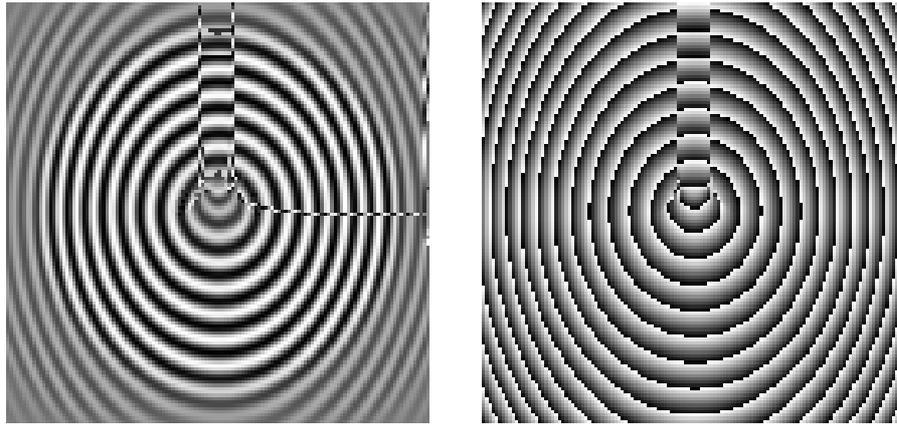


figure 5: "elliptic example": reconstruction (a) and phase of the skeleton (b)

In such a case, the reconstruction from the skeleton of the wavelet transform (not presented here) is less close to the original image than previously. This is not surprising, because such an image can hardly be modelled in the form $Ae^{i\phi}$ with slowly varying A and $\nabla\phi$ functions (it is in some sense a limit image, because it is coded with only one bit per pixel). Nevertheless, it is still possible to run the same algorithm to catch the other harmonics of the image (i.e. taking as initial guess for the a -parameter successively twice, three times ... the previous value), and to reconstruct an approximation of the image as a sum over the harmonics.

Notice that in such an example (taken from [3]), since the local frequencies are everywhere the same ones (except at the right top corner of the image), the local frequency information is of course not sufficient to segment the image (except the above mentioned corner). For the segmentation problem, additional information (such as that studied in [3]) is needed.

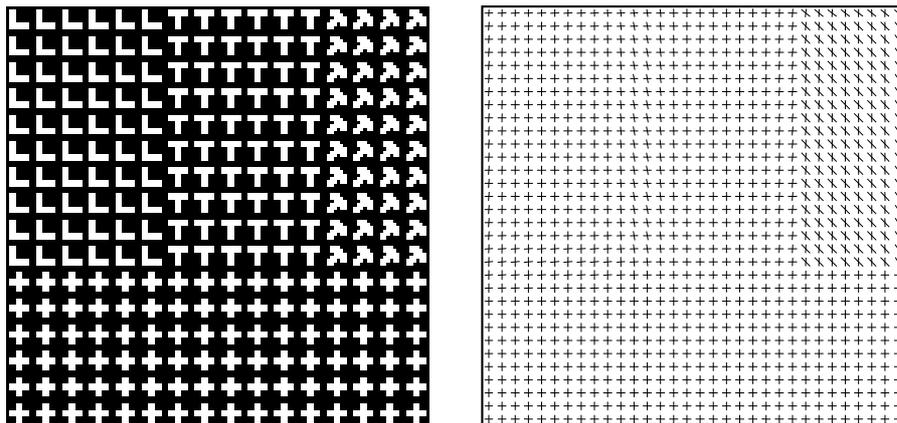


figure 6: regular texture: source image (a) and ridge (b)

Finally, it can be seen that the algorithm is not too perturbed by the presence of

noise. This can be heuristically explained as follows.

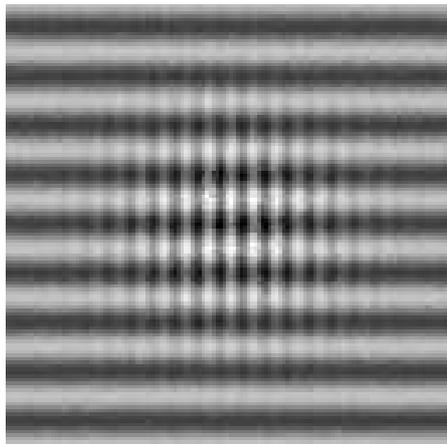


figure 7: two harmonic components in the presence of jitter noise

A given wavelet coefficient $T_s(\underline{b}, a, \theta)$ will pick only the component of noise near the point $(\underline{b}; a^{-1} \underline{r}_\theta \cdot \underline{k}_0)$ in the phase space. If the contribution of the noise near that point is much smaller than that of the feature of the image we are interested in, the approximation the algorithm is based on will still be relevant, and the answer given by the method will be a good approximation. This is illustrated in figure 7, in which a simple image formed by the superposition of two harmonic components, one of which is amplitude-modulated by a gaussian function, has been corrupted by a ± 2 bits jitter noise. Figure 8 shows the reconstruction of the two components independently, which is of quite good quality (except near the boundaries of the second component, where one can see the effect of the discontinuity at the boundary of the image). The same kind of remark applies in the case of additive gaussian noise.

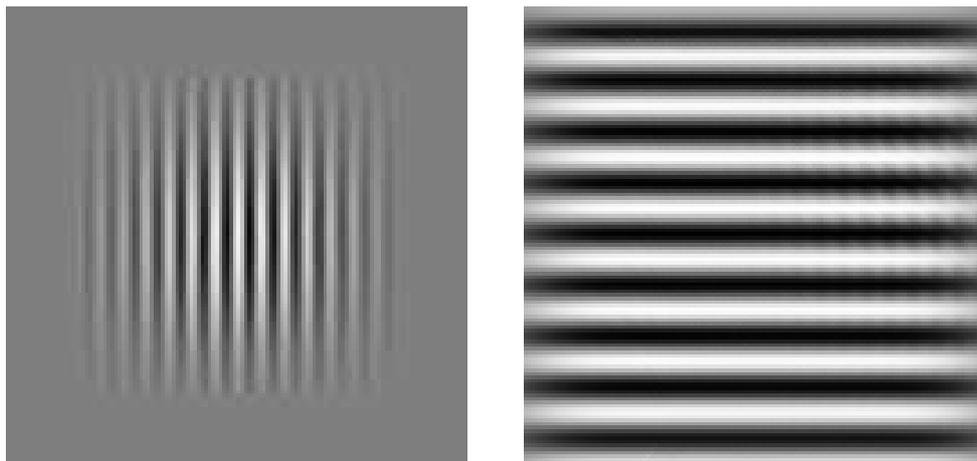


figure 8: two harmonic components in the presence of jitter noise: reconstruction

V. CONCLUSIONS.

We have described in this paper an algorithm for the analysis and characterization of local frequencies in two-dimensional signals and images (notice that using a three-dimensional wavelet transform with rotations, the algorithm can be generalized to the 3D situation). The algorithm is based on a careful analysis of the phase of the wavelet transform.

The starting point of our analysis (i.e. the study of the phase of the response of position and frequency localized filters) is similar to that of methods discussed in the literature (see [3], [2] for instance). The main difference is that our analysis is based on approximate expressions for *continuous wavelet transform*, that lead to

- Fast fixed-point algorithm for the characterization of amplitude and frequency modulated components in the original signal.
- Approximate reconstruction formulas for the extracted components.

The algorithm gives very good results whenever the frequency content of the analyzed image is not too rich, otherwise stated when the image can be well approximated as a superposition of components of the form $Ae^{i\phi}$, with a and $\nabla\phi$ slowly varying, and when the different components are localized in different regions of the phase space. In such a case, running the algorithm with appropriate initial parameters (and appropriate localization parameters for the analyzing wavelet) will yield approximate values for all the local wavevector fields. When these assumptions are not fulfilled, the precision of the approximations is smaller, and the local-frequency estimations as well as the reconstructions are biased. A careful analysis of the error (along the lines of [4]) is needed.

Nevertheless, the method described in this paper should be useful in a wide range of contexts, where such problems are not present.

FIGURE CAPTIONS

Figure 1: Localization of wavelets in the Fourier space.

Figure 2: Sum of two "chirped" components: source image (a) and ridge (b).

Figure 3: Reconstruction of the two components from the skeleton.

Figure 4: "Elliptic example": source image (a) and ridge (b).

Figure 5: "Elliptic example": reconstruction (a) and phase of the skeleton (b).

Figure 6: Regular texture: source image (a) and ridge (b).

Figure 7: Two harmonic components in the presence of jitter noise.

Figure 8: Two harmonic components in the presence of jitter noise: reconstruction.

REFERENCES.

- [1]: J.P. Antoine, P. Carrette, R. Murenzi, B. Piette, "Image analysis with 2-dimensional wavelet transform", Signal Processing 31 (1993), to appear.
- [2]: A.C. Bovik, N. Gopal, T. Emmoth, A. Restrepo (Palacios), "Localized measurements of emergent frequencies by Gabor wavelets", IEEE Trans. Inf. Th. 38 (1992) 691-712.
- [3]: J.M.H. du Buf, P. Heitkamper, "Texture features based on Gabor phase", Signal Processing 23 (1991) 227-244.
- [4]: N. Delprat, Analyse temps-fréquence de sons musicaux: exploration d'une nouvelle méthode d'extraction de données pertinentes pour un modèle de synthèse, PhD Thesis, LMA Marseille (1992).
- [5]: N. Delprat, B. Escudié, P. Guillemain, R. Kronland-Martinet, Ph. Tchamitchian, B. Torrèsani, "Asymptotic wavelet and Gabor analysis: extraction of instantaneous frequencies", IEEE Trans. Inf. Th. 38, special issue on wavelet and multiresolution analysis (1992) 644-664.
- [6]: R.B. Dingle, Asymptotic Expansions: their derivation and interpretation, Academic Press, New-York (1973).
- [7]: B. Escudié, A. Grossmann, R. Kronland-Martinet, B. Torresani, "Transformée en ondelettes de signaux asymptotiques: Emploi de la phase stationnaire", Proceedings of the GRETSI Symposium, Juan-Les-Pins (1989).
- [8]: C. Gonnet, "Caractérisation de textures à l'aide de la transformée en ondelettes", Preprint CPT-92/P.2730 (1992).
- [9]: L. Hormander, The analysis of partial differential operators I, Springer Verlag.
- [10]: J.M. Innocent, J.Y. Vinet, "Time-Frequency analysis of gravitational signals from coalescing binaries", VIRGO Note (1992).
- [11]: S. Mallat, W.L. Hwang, "Singularity detection and processing with wavelets", IEEE Trans. Inf. Th. 38 (1992), special issue on Wavelet transforms and multiresolution signal analysis, 617-643.
- [12]: R. Murenzi, Ondelettes multidimensionnelles et applications à l'analyse d'images, Thesis, IPT Louvain La Neuve, Belgium (1990).
- [13]: K. Seip, "Some remarks on a method for detecting spectral lines in a signal", Preprint CPT-89/P.2252, Marseille (1989).

APPENDIX A: APPROXIMATE EXPRESSIONS.

A.1. Stationary phase approximation.

We just describe here the stationary phase method in two-dimensions, referring to [6] for example for more details, and for the computation of all the terms of the series.

To estimate the integral

$$I = \int_{\mathbb{R}^2} M(\underline{x}) e^{i\Omega(\underline{x})} d\underline{x}$$

assuming the existence of a unique stationary point \underline{x}_0 , $\Omega(\underline{x}) - \Omega(\underline{x}_0)$ has a quadratic behaviour around $\underline{x} \cong \underline{x}_0$:

$$i(\Omega(\underline{x}) - \Omega(\underline{x}_0)) \cong \frac{i}{2}(\underline{x} - \underline{x}_0) \cdot \underline{H} \cdot (\underline{x} - \underline{x}_0)$$

where \underline{H} is the matrix of second derivatives of Ω at \underline{x}_0 . The change of variables $-\underline{u}^2 = i(\Omega(\underline{x}) - \Omega(\underline{x}_0))$, together with a Taylor series expansion of $M[\underline{x}(\underline{u})]$ around $\underline{u} = 0$ thus yields

$$I = \frac{2\pi}{\sqrt{\det H}} M(\underline{x}_0) e^{i\Omega(\underline{x}_0)} + \text{remainder}$$

The computation of the Jacobian $J(\underline{u})$ of the change of variable $\underline{x} \rightarrow \underline{u}$ leads to:

$$\begin{cases} I = I_0 + \text{remainder} \\ I_0 = \frac{2\pi}{\sqrt{|\Omega_{20}\Omega_{02} - \Omega_{11}^2|}} M(\underline{x}_0) e^{i\Omega(\underline{x}_0)} e^{-i\frac{\pi}{2} \text{sgn}[\Omega_{20}\Omega_{02} - \Omega_{11}^2]} \end{cases}$$

where we have set $\Omega_{ij} = \partial_i \partial_j \Omega$. Setting $M(\underline{x}) = A(\underline{x}) A_\psi(\underline{r}_{-\theta} \cdot \frac{\underline{x}-\underline{b}}{a})$ and $\Omega(\underline{x}) = \Phi_{(\underline{b}, a, \theta)}(\underline{x})$ yields the desired expressions.

Warning: For such an expression to be valid, it is necessary to assume that the determinant of \underline{H} doesn't vanish, i.e. that $\Omega_{20}\Omega_{02} - \Omega_{11}^2 \neq 0$. If the determinant is zero, it is then necessary to use another approximation, involving a higher order development of the phase of the integrand.

Remark: The remainder can be estimated as follows. The contribution of the second term in the Taylor series development of $M(\underline{x}(\underline{u}))/J(\underline{u})$ vanishes (because it is proportional to $\int \underline{u} \exp\{-u^2\} d\underline{u}$), and then it is necessary to go to the next term. One finally gets

$$|T_s(\underline{b}, a, \theta) - T_s(\underline{b}, a, \theta)_0| \leq K \max_{i,j} \sup_{\underline{x} \in aI_\psi + \underline{b}} \left| \partial_{ij}^2 \frac{M(\underline{x}(\underline{u}))}{J(\underline{u})} \right|$$

A.2. Wavelets of Gaussian amplitude.

The stationary phase calculus being based on a Gaussian integral computation, it can be slightly modified, by using an analyzing wavelet $\psi(\underline{x})$ with a Gaussian amplitude

$A(\underline{x}) = \exp \{ - \underline{x} \cdot \underline{A} \cdot \underline{x} \}$ where \underline{A} is a positive-definite symmetric matrix. We then have to evaluate an integral of the form

$$T_s(\underline{b}, a, \theta) = \frac{1}{a^2} e^{i\Phi(\underline{b}, a, \theta)(\underline{x}_s)} \int_{\mathbb{R}^2} A(\underline{x}) e^{-\left(\underline{r}_{-\theta} \cdot \frac{\underline{x}-\underline{b}}{a}\right) \cdot \underline{A} \left(\underline{r}_{-\theta} \cdot \frac{\underline{x}-\underline{b}}{a}\right)} e^{i\{\Phi(\underline{b}, a, \theta)(\underline{x}) - \Phi(\underline{b}, a, \theta)(\underline{x}_s)\}} d\underline{x}$$

Now, let us make the brutal first order approximation

$$\begin{cases} A(\underline{x}) \cong A(\underline{x}_s) , \\ \Phi(\underline{b}, a, \theta)(\underline{x}) - \Phi(\underline{b}, a, \theta)(\underline{x}_s) \cong \frac{1}{2}(\underline{x} - \underline{x}_s) \cdot H \cdot (\underline{x} - \underline{x}_s) . \end{cases}$$

Set $M_\theta = \frac{1}{a^2} \underline{r}_\theta M \underline{r}_{-\theta}$. This yields the approximate expression

$$\begin{aligned} T_f(\underline{b}, a, \theta) &= \frac{1}{a^2} A(\underline{x}_s) e^{i\Phi(\underline{b}, a, \theta)(\underline{x}_s)} e^{-(\underline{x}_s - \underline{b}) \cdot M_\theta \cdot (\underline{x}_s - \underline{b})} \\ &\times \int_{\mathbb{R}^2} e^{-(\underline{x} - \underline{x}_s) \cdot (M_\theta - \frac{i}{2} H) (\underline{x} - \underline{x}_s)} e^{-2(\underline{x} - \underline{x}_s) \cdot M_\theta (\underline{x}_s - \underline{b})} d\underline{x} + R(\underline{b}, a, \theta) \end{aligned}$$

Setting $\Delta = M_\theta - \frac{i}{2} H$, an inversible symmetric matrix, and $U = \underline{x} - \underline{x}_s$, and factorizing

$$(U + V) \cdot \Delta \cdot (U + V) = U \cdot \Delta \cdot U + 2U \cdot \Delta \cdot V + V \cdot \Delta \cdot V$$

with $V = \Delta^{-1} M_\theta (\underline{x}_s - \underline{b})$, we are left with a Gaussian integral, that can be easily computed. As a result, we finally have

$$T_f(\underline{b}, a, \theta) = \frac{\pi}{a^2 \sqrt{|\det \Delta|}} A(\underline{x}_s) A_\psi \left(\underline{r}_{-\theta} \cdot \frac{\underline{x} - \underline{b}}{a} \right) e^{i\Phi(\underline{b}, a, \theta)(\underline{x}_s)} e^{V \cdot \Delta \cdot V} e^{-\frac{i}{2} \text{Arg}[\det \Delta]}$$

where:

$$\begin{cases} \Delta = \underline{r}_\theta M \underline{r}_{-\theta} - \frac{i}{2} H \\ V = \Delta^{-1} \underline{r}_\theta M \underline{r}_{-\theta} (\underline{x}_s - \underline{b}) \end{cases}$$

Such expressions are easily evaluated in the case of a chirp wavelet

$$\psi(x, y) = e^{-(x^2 + y^2)} e^{ik_0(x + \alpha x^2/2)}$$

or a Morlet wavelet ($\alpha = 0$).

APPENDIX B: EVALUATION OF THE CORRECTIVE FUNCTION ON THE RIDGE.

We will restrict our analysis to the approximation described in the first part of the appendix A, i.e. to the approximate expression obtained by direct application of the

stationary phase approximation. We then need the second order derivatives of the phase function

$$\Phi_{(\underline{b}, a, \theta)}(\underline{x}) = \phi(\underline{x}) - \phi_\psi\left(\underline{r}_{-\theta} \cdot \frac{\underline{x} - \underline{b}}{a}\right)$$

evaluated on the ridge of the wavelet transform, i.e. for $\underline{x}_s = \underline{b}$. Taking into account the definition of the ridge of the wavelet transform, the first derivative reads:

$$\nabla \cdot \Phi = a_r^{-1} \underline{r}_{\theta_r} \cdot (\nabla \cdot \phi_\psi)(0) - a^{-1} \underline{r}_{-\theta} \cdot (\nabla \cdot \phi_\psi)\left(\underline{r}_{-\theta} \cdot \frac{\underline{x} - \underline{b}}{a}\right)$$

Differentiating once more, and restricting to the ridge yields

$$\left\{ \begin{array}{l} \partial_{11}^2 \Phi = -a_r^{-2} (\partial_1 a_r) [\cos \theta_r (\partial_1 \phi_\psi)(0) - \sin \theta_r (\partial_2 \phi_\psi)(0)] \\ \quad - a_r^{-1} (\partial_1 \theta_r) [\sin \theta_r (\partial_1 \phi_\psi)(0) + \cos \theta_r (\partial_2 \phi_\psi)(0)] \\ \quad - a^{-2} [\cos^2 \theta (\partial_{11}^2 \phi_\psi)(0) + \sin^2 \theta (\partial_{22}^2 \phi_\psi)(0) \\ \quad - 2 \cos \theta \sin \theta (\partial_{12}^2 \phi_\psi)(0)] \\ \partial_{12}^2 \Phi = -a_r^{-2} (\partial_2 a_r) [\cos \theta_r (\partial_1 \phi_\psi)(0) - \sin \theta_r (\partial_2 \phi_\psi)(0)] \\ \quad - a_r^{-1} (\partial_2 \theta_r) [\sin \theta_r (\partial_1 \phi_\psi)(0) + \cos \theta_r (\partial_2 \phi_\psi)(0)] \\ \quad - a^{-2} [2 \cos \theta \sin \theta ((\partial_{11}^2 \phi_\psi)(0) - (\partial_{22}^2 \phi_\psi)(0)) \\ \quad + (\cos^2 \theta - \sin^2 \theta) (\partial_{12}^2 \phi_\psi)(0)] \\ \partial_{22}^2 \Phi = -a_r^{-2} (\partial_2 a_r) [\sin \theta_r (\partial_1 \phi_\psi)(0) - \cos \theta_r (\partial_2 \phi_\psi)(0)] \\ \quad - a_r^{-1} (\partial_2 \theta_r) [\cos \theta_r (\partial_1 \phi_\psi)(0) - \sin \theta_r (\partial_2 \phi_\psi)(0)] \\ \quad - a^{-2} [\sin^2 \theta (\partial_{11}^2 \phi_\psi)(0) + \cos^2 \theta (\partial_{22}^2 \phi_\psi)(0) \\ \quad + 2 \cos \theta \sin \theta (\partial_{12}^2 \phi_\psi)(0)] \end{array} \right.$$

In the case of a "chirp wavelet"

$$\psi(x, y) = e^{-(x^2+y^2)} e^{ik_0(x+\alpha x^2/2)}$$

one gets

$$\left\{ \begin{array}{l} \partial_{11}^2 \Phi = -\frac{k_0}{a_r^2} \cos \theta_r \partial_1 a_r - \frac{k_0}{a_r} \sin \theta_r \partial_1 \theta_r - \frac{k_0 \alpha}{a_r^2} \cos^2 \theta_r \\ \partial_{12}^2 \Phi = -\frac{k_0}{a_r^2} \sin \theta_r \partial_2 a_r + \frac{k_0}{a_r} \cos \theta_r \partial_2 \theta_r - \frac{k_0 \alpha}{a_r^2} \sin^2 \theta_r \\ \partial_{22}^2 \Phi = -\frac{k_0}{a_r^2} \cos \theta_r \partial_2 a_r - \frac{k_0}{a_r} \sin \theta_r \partial_2 \theta_r - \frac{k_0 \alpha}{a_r^2} \cos \theta_r \sin \theta_r \end{array} \right.$$

and the case of the Morlet wavelet

$$\psi(x, y) = e^{-(x^2+y^2)} e^{ik_0 x}$$

is obtained by setting $\alpha = 0$ in the previous formula.