

Fast Matching Pursuit with a multiscale dictionary of Gaussian chirps

Rémi Gribonval

Abstract— We introduce a modified Matching Pursuit algorithm, called Fast Ridge Pursuit, to approximate N -dimensional signals with M Gaussian chirps at a computational cost $\mathcal{O}(MN)$ instead of the expected $\mathcal{O}(MN^2 \log N)$. At each iteration of the pursuit, the best Gabor atom is first selected, then its scale and chirp rate are locally optimized so as to get a “good” chirp atom, *i.e.* one for which the correlation with the residual is locally maximized. A ridge theorem of the Gaussian chirp dictionary is proved, from which an estimate of the locally optimal scale and chirp is built. The procedure is restricted to a sub-dictionary of local maxima of the Gaussian Gabor dictionary, so as to accelerate further the pursuit. The efficiency and speed of the method is demonstrated on a sound signal.

Keywords— adaptive signal processing, signal representations, approximation methods, time-frequency analysis, chirp modulation, frequency estimation, redundant systems, complexity theory.

I. INTRODUCTION

There has been a considerable interest last decade in developing analysis techniques to decompose non-stationary signals into elementary components, called *atoms*, that characterize their salient features. As many signals display both oscillatory phenomena, which time-frequency methods can extract, and transients or singularities, to which time-scale techniques [1], [2], [3] are better adapted [4], [5], [6], adaptive decompositions were developed, using *redundant* families of atoms that can characterize independently scale and frequency (Local Cosine [7], Wavelet Packets [8], Gabor multiscale dictionary [9], [10]).

Chirp atoms were introduced to deal with the non-stationary behavior of the instantaneous frequency of some signals [11]. Baraniuk and Jones [12] built orthonormal bases and frames of such chirp atoms, while Mann and Haykin [13] defined a “Chirplet Transform”. Roughly speaking, this transform compares a signal $x(t)$ with each chirp atom

$$g_{(s,u,\xi,c)}(t) = \frac{1}{\sqrt{s}} g\left(\frac{t-u}{s}\right) e^{i(\xi(t-u) + \frac{c}{2}(t-u)^2)} \quad (1)$$

of a large family, the *chirp dictionary* \mathcal{D}^+ , which is an extension of the Gabor multiscale time-frequency dictionary \mathcal{D} [9] [10]. These atoms are characterized by their scale s ,

R. Gribonval is with IRISA-INRIA, a facility of the French National Center for Computer Science and Control (INRIA). Part of this paper was written during the author’s graduate studies at the Centre de Mathématiques Appliquées (CMAP), École Polytechnique, France, and his postdoctoral year in the Industrial Mathematics Institute (IMI) of the Department of Mathematics of the University of South Carolina, with support from the National Science Foundation (NSF) Grant No. DMS-9872890. Address: IRISA-INRIA, Campus de Beaulieu, 35042 Rennes cedex, France, Phone: +33 (0) 2 99 84 25 06, Fax: +33 (0) 2 99 84 25 32, E-mail: Remi.Gribonval@inria.fr

time u , frequency ξ and *chirp rate* c . Their instantaneous frequency $\omega(t) = \xi + c(t-u)$ varies linearly with time.

In an orthonormal basis of chirp atoms [12], a given signal can be efficiently decomposed into elementary chirps. However the elementary atoms are somehow too “rigid” for many applications, as their parameters s , $\xi \propto 1/s$ and $c \propto 1/s^2$ are not independent one from another. On the other hand, the chirplet transform is very redundant and does not have this intrinsic rigidity. It can thus provide with a large variety of viewpoints to look at the signal, in order to find meaningful structures in it. However its redundancy is also its weakness as it makes the computational complexity of the chirplet transform very large.

Bultan [14] suggested to use the Matching Pursuit algorithm of Mallat and Zhang [15] to decompose a signal into elementary chirp atoms. He demonstrated the interest of this technique, but its practical use was limited by the large computational complexity $\mathcal{O}(MN^2 \log N)$ needed to get an M -term approximation of an N -sample signal. In order to limit the complexity, Bultan suggested to reduce the size of the dictionary by limiting the resolution of the chirp rate.

In this work we show that it is possible to get rid of such a limitation *and* get a low complexity $\mathcal{O}(MN)$ by modifying the underlying “Matching Pursuit” algorithm and using a *Gaussian* chirp dictionary. To get such a low complexity we introduce a (substantially) modified pursuit algorithm, by using some *ridge* techniques and the *local maxima* of the Gabor dictionary.

The paper is organized as follows. In the next section, we review the definition of the multiscale time-frequency chirp dictionary \mathcal{D}^+ and show the numerical complexity implied by its very large size. In section III the definition and basic properties of the Matching Pursuit are recalled. Section IV is devoted to the detailed study of the ridges of the Gaussian multiscale Gabor dictionary. We use those results to analyze the selection of the *locally optimal chirp atom*. In section V we summarize the Ridge Pursuit algorithm with the real-valued chirp dictionary, and show how it can be further accelerated with a sub-dictionary technique. Finally we analyze in section VI the numerical results obtained with our new algorithm on an acoustic signal.

II. MULTISCALE DICTIONARY OF TIME-FREQUENCY CHIRP ATOMS

Every chirp atom (1) is obtained from an elementary window $g(t)$ by dilation, translation, frequency and chirp modulation. It can thus be described with its index (s, u, ξ, c) . The window $g(t)$ is localized around 0 both in the time domain and the frequency domain. As a re-

sult $g_{(s,u,\xi,c)}$ is localized at time u with a temporal dispersion proportional to its scale s . The Wigner-Ville distribution $WV[g_{(s,u,\xi,c)}](t,\omega)$ [16], [17] of a chirp atom defines a quadratic time-frequency energy distribution. It is localized around the line of instantaneous frequency $\omega(t) = \xi + c(t-u)$. Its dispersion is proportional to $1/s$ in the ω direction. A Gaussian chirp atom is built from the unit Gaussian window $g(t) = \pi^{-1/4}e^{-t^2/2}$. Such an atom is displayed on Figure 1 with its Wigner-Ville distribution.

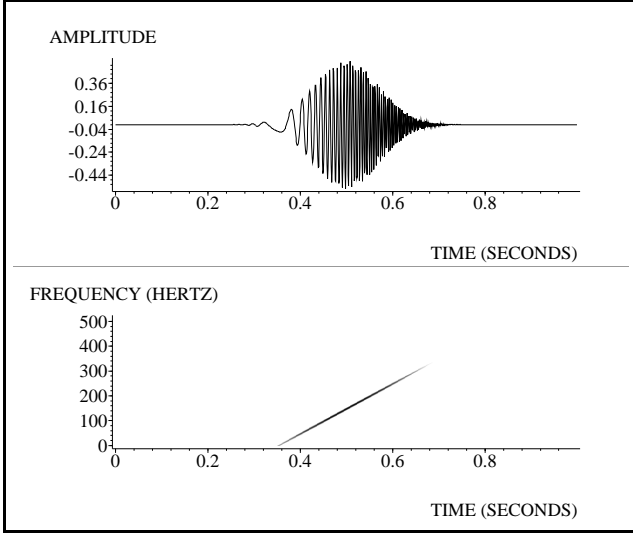


Fig. 1. A Gaussian chirp atom (top) and its Wigner-Ville distribution (bottom). The energy density is grey-coded from the smallest values (white) to the largest values (black).

A. Sampling the dictionary

The set $\{g_{(s,u,\xi,0)}, (s,u,\xi) \in \mathbb{R}^+ \times \mathbb{R} \times \mathbb{R}\}$ of chirp atoms with chirp rate $c = 0$ is exactly the multiscale Gabor dictionary [9], [15], [10]. The discrete Gabor dictionary \mathcal{D} is the collection of atoms $g_{(s,u,\xi,0)}$ (denoted for short by $g_{(s,u,\xi)}$) such that $(s,u,\xi) = (a^j, na^j \Delta u, ka^{-j} \Delta \xi)$, $j, n, k \in \mathbb{Z}$, where Δu and $\Delta \xi$ are some constants. Watson and Gilholm [18] showed that this sampling of the scale, time and frequency parameters is uniform with respect to the natural Riemannian metric of the continuous dictionary induced by $d(g_{\gamma_1}, g_{\gamma_2}) = 1 - |\langle g_{\gamma_1}, g_{\gamma_2} \rangle|$, where $\langle x(t), y(t) \rangle = \int_{-\infty}^{+\infty} x(t)\overline{y(t)}dt$ is the standard inner product on $L^2(\mathbb{R})$. The same point of view leads to sampling the chirp rate as $c = la^{-2j} \Delta c$, $l \in \mathbb{Z}$. The discrete chirp dictionary \mathcal{D}^+ is thus the family of atoms $g_{(s,u,\xi,c)}$ such that $(s,u,\xi,c) \in \Gamma^+$, where

$$\Gamma^+ = \{(a^j, na^j \Delta u, ka^{-j} \Delta \xi, la^{-2j} \Delta c), j, n, k, l \in \mathbb{Z}\}. \quad (2)$$

As the set of atoms at a given scale j and chirp rate l is a Weyl-Heisenberg frame, it can only span $L^2(\mathbb{R})$ if $\Delta u \times \Delta \xi \leq 2\pi$ [19]. When $\Delta u \times \Delta \xi < 2\pi$, \mathcal{D} is complete [15], thus $\mathcal{D}^+ \supset \mathcal{D}$ is also complete.

B. Size of the discrete chirp dictionary

The size of the discrete chirp dictionary \mathcal{D}^+ is a function of the sampling steps $a, \Delta u, \Delta \xi$ and Δc . When analyzing

a discrete N point signal, one also has to consider the limitations of the sampling rate and the signal size. The scale a^j can thus only vary between 1 and N , which makes a total of $\mathcal{O}(\log N)$ scales. At each scale, there are $\mathcal{O}(N)$ sampled values of $(u, \xi) \in [0, N-1] \times [-\pi, \pi]$. Because of the Nyquist condition, the instantaneous frequency is constrained to $\xi + c(t-u) \in [-\pi, \pi]$, $\forall t \in [u - a^j/2, u + a^j/2]$, i.e. $|c|a^j/2 \leq \pi - |\xi|$. For given a^j and ξ , the chirp rate c can take $\mathcal{O}((\pi - |\xi|)a^j)$ values. On the average, at scale a^j , it thus takes $\mathcal{O}(a^j)$ distinct values. The total number of chirp atoms in the discrete chirp dictionary \mathcal{D}^+ is thus of the order of $\sum_{j=0}^{\log N} \mathcal{O}(N/a^j) \times \mathcal{O}(a^j) \times \mathcal{O}(a^j) = \mathcal{O}(N^2)$.

III. STANDARD MATCHING PURSUIT WITH \mathcal{D}^+

The Matching Pursuit [15] is a greedy strategy to decompose a signal x into a linear combination of atoms chosen among a dictionary $\mathcal{D} = \{g_\gamma, \gamma \in \Gamma\}$, i.e. a redundant family of unit vectors in a Hilbert space \mathcal{H} . It iteratively defines an m^{th} -order residual $R^{m-1}x$ (starting with $R^0x = x$) in the following way :

1. Compute $|\langle R^{m-1}x, g_\gamma \rangle|^2$ for all $\gamma \in \Gamma$.
2. Select the best atom of the dictionary

$$g_{\gamma_m} = \arg \max_{\gamma \in \Gamma} |\langle R^{m-1}x, g_\gamma \rangle|^2. \quad (3)$$

3. Compute the new residual by removing the component along the selected atom

$$R^m x = R^{m-1}x - \langle R^{m-1}x, g_{\gamma_m} \rangle g_{\gamma_m}. \quad (4)$$

After M iterations, one gets an M -term approximation $x_M = x - R^M x = \sum_{m=1}^M \langle R^{m-1}x, g_{\gamma_m} \rangle g_{\gamma_m}$. The energy is split among the selected components as $\|x\|^2 = \sum_{m=1}^M |\langle R^{m-1}x, g_{\gamma_m} \rangle|^2 + \|R^M x\|^2$. The Matching Pursuit is very similar to the Projection Pursuit principle introduced in statistics by Huber [20], which strong convergence $\lim_{M \rightarrow \infty} \|R^M x\| = 0$ was proved by Jones [21] whenever the dictionary \mathcal{D} is complete, i.e. $\overline{\text{span}}(\mathcal{D}) = \mathcal{H}$.

Let us note that the Matching Pursuit does *not* provide the best approximation to x by a linear combination of M atoms from \mathcal{D} . Actually, getting such a best M -term approximant is an NP-hard problem [22]. In finite dimension N , at most N atoms should be needed to represent a signal x , but in general the Matching Pursuit goes on forever without ever giving an exact decomposition. This can be fixed with a variant, the Orthonormal Matching Pursuit [23]. However, as the Orthonormal Matching Pursuit performs a Gram-Schmidt orthonormalization of the family $\{g_m\}_{m \geq 1}$, its computational cost is significantly higher than that of the “pure” Matching Pursuit.

With the chirp dictionary \mathcal{D}^+ and a N -point signal, the computation of $\langle R^{m-1}x, g_\gamma \rangle, g_\gamma \in \mathcal{D}^+$ can be done with $\mathcal{O}(N^2 \log N)$ operations, using FFT-based algorithms with appropriate windows [14], [18]. The search for the “best” atom (3) costs $\mathcal{O}(N^2)$, and the update of the residual (4) only costs $\mathcal{O}(N)$, hence the total complexity $\mathcal{O}(MN^2 \log N)$ of M iterations of pursuit with the chirp dictionary. Such a “brute force” Chirp Matching Pursuit

is thus limited to the analysis of small signals, with only few iterations.

IV. Ridge Pursuit

Because of the large size of \mathcal{D}^+ , one cannot afford to compute the correlation $\langle R^{m-1}x, g_{(s,u,\xi,c)} \rangle$ of the residual with *every* atom of \mathcal{D}^+ . As a consequence, the choice of the “best” atom $g_{(s_m, u_m, \xi_m, c_m)}$ must be done in an *approximate* way. In other words, one needs to “guess” where a “good” chirp atom is located, without scanning the whole dictionary.

One can notice that the chirp dictionary \mathcal{D}^+ is only an extension of the Gabor dictionary \mathcal{D} . As \mathcal{D} is complete, the set of inner products $\langle R^{m-1}x, g_\gamma \rangle$, $g_\gamma \in \mathcal{D}$ contains all the information available about $R^{m-1}x$. It is thus theoretically sufficient to compute these inner products to select the best chirp atom. We will actually show, with Theorem 1, that the behavior of $\langle R^{m-1}x, g_\gamma \rangle$ in the neighborhood of the best Gabor atom contains enough information to select a “locally optimal” chirp atom. A “good” chirp atom $g_{(s_m, u_m, \xi_m, c_m)}$ is selected with a two-step pursuit. First, one selects the best Gabor atom

$$g_{(s_m^*, u_m^*, \xi_m^*)} \triangleq \arg \max_{g_{(s,u,\xi)} \in \mathcal{D}} |\langle R^{m-1}x, g_{(s,u,\xi)} \rangle|. \quad (5)$$

Then one explores its neighborhood in \mathcal{D}^+ , so as to find a good chirp atom

$$g_{(s_m, u_m, \xi_m, c_m)} \triangleq \arg \max_{g_{(s, u_m^*, \xi_m^*, c)}} |\langle R^{m-1}x, g_{(s, u_m^*, \xi_m^*, c)} \rangle| \quad (6)$$

by selecting locally optimal chirp rate and scale parameters s_m and c_m . The time and frequency parameters u_m^* and ξ_m^* are kept constant. Generally speaking, we could allow for re-optimization of the time and frequency parameters too. However we chose not to re-estimate them, because the re-optimized values are very close to the initial ones in practice. On the contrary, the re-optimized values of s_m and c_m can be substantially different from the initial ones.

One can see that after selecting the best Gabor atom (5), the second step (6) implies an *exhaustive scanning* of the neighborhood of this atom. However this scanning is still very costly. We replace it by a *fast estimation* $\mathcal{O}(1)$ of s_m and c_m , using again Theorem 1, which helps us extract the information we need from the local behavior of $(s, u, \xi, c) \mapsto \langle R^{m-1}x, g_{(s,u,\xi,c)} \rangle$ in the neighborhood of the best Gabor atom. We hereby define a *Ridge Pursuit*, which complexity $\mathcal{O}(MN \log^2 N)$ is identical to that of the standard Matching Pursuit with the Gabor dictionary \mathcal{D} . Let us outline one step of the Ridge Pursuit.

1. Select the best Gabor atom $g_{(s_m^*, u_m^*, \xi_m^*)}$.
2. Use the local behavior of $\xi \mapsto |\langle R^{m-1}x, g_{(s_m^*, u_m^*, \xi)} \rangle|$ in the neighborhood of ξ_m^* to estimate the chirp parameter c_m and get a better estimate of the scale parameter s_m .
3. Compute the new residual using the chirp atom $g_{(s_m, u_m^*, \xi_m^*, c_m)}$.

A. Ridges of the Gaussian chirp dictionary

Discrete signals $x_d \in \mathbb{R}^N$ are obtained by sampling band-limited continuous-time signals $x(t)$, and the discrete inner products $\langle x_d, y_d \rangle = \sum_{n=0}^{N-1} x_d[n] \overline{y_d[n]}$ are close to their continuous counterparts $\langle x, y \rangle = \int_{-\infty}^{+\infty} x(t) \overline{y(t)} dt$. Chirplets are most useful for the representation of signals that contain well defined instantaneous frequency $\omega(t)$ and chirp-rate $\omega'(t)$. From now on, we consider the model $R^{m-1}x(t) = a(t)e^{i\phi(t)}$ where these quantities $\omega(t) = \phi'(t)$ and $c(t) = \phi''(t)$ are easily defined. Our results can be extended to the case of a superposition of finitely many such continuous signals, provided a sufficient separation of their instantaneous frequencies is granted.

The goal of the following ridge theorem (which is proved in appendix A) is to show that, under certain regularity conditions, the residual $R^{m-1}x$, seen “through” a Gaussian chirp atom $g_{(s,u,\xi,c)}$, looks like another Gaussian chirp atom $g_{\gamma^+(u)} \in \mathcal{D}^+$, i.e. $\langle R^{m-1}x, g_{(s,u,\xi,c)} \rangle \approx A(u)e^{i\Phi(u)} \langle g_{\gamma^+(u)}, g_{(s,u,\xi,c)} \rangle$.

Theorem 1: Let $R^{m-1}x(t) = a(t)e^{i\phi(t)}$. Suppose that $\|a\|_\infty < \infty$, $\|\phi'''\|_\infty < \infty$ and $\|b'''\|_\infty < \infty$, with $b(t) \triangleq -\log a(t)$. Let u a time where $b'(u) > 0$, and let $g_{(s,u,\xi,c)}$ a Gaussian chirp atom. Then

$$\langle R^{m-1}x, g_{(s,u,\xi,c)} \rangle = A(u)e^{i\Phi(u)} (\langle g_{\gamma^+(u)}, g_{(s,u,\xi,c)} \rangle + \epsilon(s, u, \xi, c)) \quad (7)$$

where

$$A(u) = \frac{a(u)}{(b''(u)/\pi)^{\frac{1}{4}}} e^{\frac{(b'(u))^2}{2b''(u)}} \quad (8)$$

$$\Phi(u) = \phi(u) - \phi'(u) \frac{b'(u)}{b''(u)} + \frac{\phi''(u)}{2} \left(\frac{b'(u)}{b''(u)} \right)^2 \quad (9)$$

$$\gamma^+(u) = \left(\frac{1}{\sqrt{b''(u)}}, u - \frac{b'(u)}{b''(u)}, \phi'(u) - \phi''(u) \frac{b'(u)}{b''(u)}, \phi''(u) \right) \quad (10)$$

and $|\epsilon(s, u, \xi, c)|$ is bounded by

$$\epsilon_{max}(s, u) = \frac{(b''(u))^{\frac{1}{4}} s^{\frac{1}{2}}}{\pi^{\frac{1}{2}}} \left(\frac{K s^3 \sigma_3^3}{6} e^{\frac{1}{6}} + \frac{4 \|a\|_\infty}{a(u)} \frac{(K s^3)^{\frac{1}{3}}}{\exp \frac{1}{2(K s^3)^{2/3}}} \right) \quad (11)$$

with $K \triangleq \|b'''\|_\infty + \|\phi'''\|_\infty$ and $\sigma_3^3 \triangleq \int |t|^3 g(t) dt$.

The hypothesis $b''(u) > 0$ simply corresponds to $a''(u)/a(u) < (a'(u)/a(u))^2$. For instance, it holds in the neighborhood of smooth local extrema of $a(u)$, where $a''(u)/a(u) \leq 0$. In particular this is the case when u is the time-location of the best Gabor atom, because $(s, u, \xi) \mapsto |\langle R^{m-1}x, g_{(s,u,\xi)} \rangle|$ is locally maximum. Moreover for such a u , $b'(u)$ is very small, hence $\gamma^+(u)$ is almost $(1/\sqrt{b''(u)}, u, \phi'(u), \phi''(u))$.

From this theorem, one can observe that if

$$\|b'''\|_\infty \ll |b''(u)|^{3/2} \quad \text{and} \quad \|\phi'''\|_\infty \ll |b''(u)|^{3/2} \quad (12)$$

then $\epsilon_{max}(1/\sqrt{b''(u)}, u) \ll 1$, so that the best chirp atom at time u is close to $g_{\gamma^+(u)}$. The locally optimal parameters

s, ξ, c can thus be obtained by estimating the index $\gamma^+(u)$. Let us now study how much information the location of the best Gabor atom gives about $\gamma^+(u)$.

B. Scale and frequency of the best Gabor atom

In the following we suppose that ϵ_{max} can be neglected. As the best Gabor atom (5) is the absolute maximum of $|\langle R^{m-1}x, g_{(s,u,\xi)} \rangle|$, it is a local maximum along s and ξ . If we additionally suppose that $b''(u) \ll 1$, then the right hand side in (7) becomes $a(u)e^{i\phi(u)} \langle e^{i(\phi'(u)(t-u) + \phi''(u)(t-u)^2/2)}, g_{(s,u,\xi)} \rangle$. As the dictionary is Gaussian, the inner product that appears in this approximant is the Fourier transform $\hat{g}_{(s,0,0,\phi''(u))}(\xi - \phi'(u))$ of a Gaussian chirp atom, which analytic expression is known [24]. For a given u , its maximum (or *ridge*) along s and ξ is located at $\xi = \phi'(u)$ and $s = 1/\sqrt{|\phi''(u)|}$. Thus one has

$$\xi_m^* \approx \phi'(u_m^*) \text{ and } s_m^* \approx 1/\sqrt{|\phi''(u_m^*)|}. \quad (13)$$

Bounds on the error of these estimates can be found in [25]. It is well known that the ridges of the wavelet transform or of the windowed Fourier transform give the instantaneous frequency [26], [17] : this result shows that the ridges of the Gabor dictionary additionally provide with the instantaneous chirp rate. Now, it is sufficient that

$$\|b'''\|_\infty \ll |\phi''(u)|^{3/2} \text{ and } \|\phi'''\|_\infty \ll |\phi''(u)|^{3/2} \quad (14)$$

to get $\epsilon_{max}(1/\sqrt{|\phi''(u)|}, u) \ll 1$ and control the location (13) of the best Gabor atom, which gives information on the locally optimal chirp rate $|c_m| \approx |\phi''(u_m^*)| \approx (s_m^*)^{-2}$. Unfortunately the estimate $\hat{c}_m = \pm (s_m^*)^{-2}$ is far from the ideal one. First, because one has to determine its sign by computing the two inner products $\langle R^{m-1}x, g_{(s_m^*, u_m^*, \xi_m^*, \pm (s_m^*)^{-2})} \rangle$. But also mainly because it is a very poor estimate when, as usual, the scale $s = 2^j$ is coarsely quantized. Thus this estimate is not sufficient to avoid the costly $\mathcal{O}(N^2)$ “scanning” of the possible chirp atoms $g_{(s, u_m^*, \xi_m^*, c)}$.

C. Fast local estimation of the best chirp atom

The *local behavior* of $\xi \mapsto \langle R^{m-1}x, g_{(s_m^*, u_m^*, \xi)} \rangle$ in the neighborhood of ξ_m^* conveys much more information about the locally optimal chirp atom $g_{\gamma^+(u)}$ than the *location* (s_m^*, u_m^*, ξ_m^*) of the best Gabor atom does. Indeed, if $\epsilon_{max}(s_m^*, u_m^*) \ll 1$, then from Theorem 1, $\langle R^{m-1}x, g_{(s_m^*, u_m^*, \xi)} \rangle \approx C \langle g_{\gamma^+(u_m^*)}, g_{(s_m^*, u_m^*, \xi)} \rangle$, where $C \in \mathbb{C}$ is some constant independent on ξ . Using the analytic expression of the inner product between two Gaussian chirp atoms [24], one can get the following spectral estimation [27], [28] of the parameters of $g_{\gamma^+(u)}$, which is proved in [25].

Proposition 1: If $\epsilon_{max}(s_m^*, u_m^*) \ll 1$, then $\langle R^{m-1}x, g_{(s_m^*, u_m^*, \xi)} \rangle \approx Ae^{i\Phi}$ where $\xi \mapsto \log A(\xi)$ and $\xi \mapsto \Phi(\xi)$ are second order polynomials in ξ with

$$\phi''(u_m^*) = -\frac{\Phi''(\xi)}{((\log A)''(\xi))^2 + (\Phi''(\xi))^2} \quad (15)$$

and

$$b''(u_m^*) + \frac{1}{(s_m^*)^2} = \frac{-\log A''(\xi)}{((\log A)''(\xi))^2 + (\Phi''(\xi))^2}. \quad (16)$$

Moreover the following bounds hold

$$|\Phi''(\xi)| \leq \frac{(s_m^*)^2}{2} \quad (17)$$

$$0 > \log A''(\xi) \geq -(s_m^*)^2. \quad (18)$$

One can easily estimate $(\log A)''(\xi)$ and $\Phi''(\xi)$ (which are independent of ξ) using only the local behavior of $\xi \mapsto \langle R^{m-1}x, g_{(s_m^*, u_m^*, \xi)} \rangle$ around the best Gabor atom. Then (17) and (18) are used to *test* the validity of the approximation $\epsilon_{max} \ll 1$. Whenever the test is negative, the Ridge Pursuit is conservative : it does not try to find a better chirp atom than the best Gabor atom, but instead keeps it as its “good chirp atom” and steps forward to the next iteration. In the case of a positive test, we will *assume* that the model is valid. Thanks to (15) and (16), the estimates of $(\log A)''(\xi)$ and $\Phi''(\xi)$ provide with estimates of $\phi''(u_m^*)$ and $b''(u_m^*)$, *i.e.* an estimate of $\gamma^+(u)$. This estimate is now obtained without costly “scanning”.

The definition of the Ridge Pursuit will be complete by showing how to efficiently estimate $(\log A)''(\xi)$ and $\Phi''(\xi)$.

D. Numerical estimation by linear interpolation

In order to get as local an estimation as possible, we estimate $(\log A)''(\xi)$ and $\Phi''(\xi)$ through a parabolic interpolation. We use three Gaussian Gabor atoms $g_\varepsilon = g_{(s_m^*, u_m^*, \xi_m^* + \varepsilon \Delta\xi / s_m^*)}$, $\varepsilon \in \{-1, 0, +1\}$, of the discrete Gabor dictionary \mathcal{D} , and their inner products $\langle R^{m-1}x, g_\varepsilon \rangle = \tilde{A}_\varepsilon e^{i\tilde{\Phi}_\varepsilon}$. These inner products were already computed for the selection of the best Gabor atom.

The numerical parabolic interpolation of $\log \tilde{A}_\varepsilon$ (resp. $\tilde{\Phi}_\varepsilon$), taking into account the frequency bin size $\Delta\xi / s_m^*$, leads to estimates

$$\Phi''(\xi) \approx (\tilde{\Phi}_{-1} - 2\tilde{\Phi}_0 + \tilde{\Phi}_1) \left(\frac{s_m^*}{\Delta\xi} \right)^2 \quad (19)$$

$$(\log A)''(\xi) \approx \log \frac{\tilde{A}_{-1}\tilde{A}_1}{\tilde{A}_0^2} \left(\frac{s_m^*}{\Delta\xi} \right)^2. \quad (20)$$

As $\tilde{\Phi}_\varepsilon$ is defined modulo 2π , the estimate of $\Phi''(\xi)$ is only defined modulo $2\pi(s_m^*/\Delta\xi)^2$. However, thanks to condition (17), its only admissible value(s) lie within the interval $[-(s_m^*)^2/2, +(s_m^*)^2/2]$. In order to eliminate the ambiguity, it is necessary and sufficient to impose that the length of this interval is strictly less than $2\pi(s_m^*/\Delta\xi)^2$, *i.e.* to choose $\Delta\xi < \sqrt{2\pi}$ in the definition of \mathcal{D} and \mathcal{D}^+ (see (2)). Thus s_m and c_m are estimated at a cost $\mathcal{O}(1)$ from the inner products $\langle R^{m-1}x, g_\varepsilon \rangle$.

V. FAST RIDGE PURSUIT

For the analysis of real-valued signal, we do not make use of complex-valued atoms (1) but of real-valued ones. They are defined [15], [14] as

$$g_{(s,u,\xi,c,\phi)} = K_{(s,u,\xi,c,\phi)} g \left(\frac{t-u}{s} \right) \cos \left(\xi(t-u) + \frac{c}{2}(t-u)^2 + \phi \right) \quad (21)$$

with some normalizing constant $K_{(s,u,\xi,c,\phi)}$. Obviously $g_{(s,u,\xi,c,\phi)}$ lies in the two-dimensional subspace $\mathcal{V}_{(s,u,\xi,c)} = \text{Span}\{g_{(s,u,\xi,c)}, \overline{g_{(s,u,\xi,c)}}\}$, and

$$\sup_{\phi} |\langle R^{m-1}x, g_{(s,u,\xi,c,\phi)} \rangle| = \|P_{\mathcal{V}_{(s,u,\xi,c)}} R^{m-1}x\|, \quad (22)$$

where $P_{\mathcal{V}}$ denotes the orthogonal projector onto the subspace \mathcal{V} . We show in appendix B that the right hand side of (22), as well as the corresponding optimal phase $\phi(s, u, \xi, c) = \arg \max_{\phi} |\langle R^{m-1}x, g_{(s,u,\xi,c,\phi)} \rangle|$, can be computed in $\mathcal{O}(1)$ from $\langle R^{m-1}x, g_{(s,u,\xi,c)} \rangle$.

Let us now summarize the Ridge Pursuit algorithm with real-valued Gaussian chirp atoms, and compute its numerical complexity. Each iteration is decomposed into a few steps.

A. Ridge Pursuit Algorithm

1. Compute $\langle R^{m-1}x, g_{(s,u,\xi)} \rangle$ for each complex Gaussian Gabor atom $[\mathcal{O}(N \log^2 N)]$.
2. Compute $\|P_{\mathcal{V}_{(s,u,\xi)}} R^{m-1}x\|$ and select the location (s_m^*, u_m^*, ξ_m^*) of the best real-valued Gaussian Gabor atom $[\mathcal{O}(N \log N)]$.
3. Estimate the locally optimal parameters s_m and c_m with a parabolic interpolation $[\mathcal{O}(1)]$.
4. Compute $\langle R^{m-1}x, g_{(s_m, u_m^*, \xi_m^*, c_m)} \rangle$, and determine the best real-valued chirp atom $g_{(s_m, u_m^*, \xi_m^*, c_m, \phi_m)}$ in $\mathcal{V}_{(s_m, u_m^*, \xi_m^*, c_m)}$ $[\mathcal{O}(N)]$.
5. Update the residual $[\mathcal{O}(N)]$.

The overall complexity of one iteration of real-valued Ridge Pursuit is $\mathcal{O}(N \log^2 N)$, hence the total cost $\mathcal{O}(MN \log^2 N)$ of M iterations. An accelerating technique was introduced by Bergeaud and Mallat [29], [30] for the Matching Pursuit analysis of images. It can be used to get a Fast Ridge Pursuit algorithm. The overall algorithm is described in full details in [25], and we give here its main features. We use *local maxima* of the Gabor dictionary \mathcal{D} , that is Gabor atoms $g_{(s,u,\xi)}$ where either $u \mapsto |\langle x, g_{(s,u,\xi)} \rangle|$ or $\xi \mapsto |\langle x, g_{(s,u,\xi)} \rangle|$ has a local maximum. A number P is fixed arbitrarily, and the following steps are done iteratively.

B. Fast Ridge Pursuit Algorithm

1. Build a sub-dictionary \mathcal{D}_m of P local maxima of the Gabor dictionary \mathcal{D} .
 2. For each atom in \mathcal{D}_m , use the fast local estimation procedure to get a good chirp atom. The collection of these chirp atoms is a sub-dictionary \mathcal{D}_m^+ of the chirp dictionary \mathcal{D}^+ .
 3. Run a “normal” pursuit in \mathcal{D}_m^+ until it is empty.
- By choosing $\log^2 N \ll P \ll N$, the overall complexity becomes $\mathcal{O}(MN)$ [25].

VI. APPLICATIONS

The Ridge Pursuit and Fast Ridge Pursuit algorithms were implemented using the Matching Pursuit Package of the LastWave program [31]. We used them to analyze a sound recording with sung voice and orchestra [32]. It

is well known that a characteristic of the sung voice is its vibrato [33], which the Gabor Matching Pursuit was not likely to decompose sparsely. The signal duration was approximately 2.5 seconds, at a sampling rate of 11,025 Hertz, so the signal length was about $N \approx 30000$ samples. A Gabor Matching Pursuit and a Fast Ridge Pursuit were computed with $M = 5000$ iterations.

One needs first to realize how high the complexity of a “brute force” Matching Pursuit with the chirp dictionary [14] would have been. With an (optimistic) average of 100 MFlops to 1 GFlops for today's computers, the $5000 \times (30000)^2 \log_2 30000 \approx 6 \times 10^{13}$ operations would have required 6×10^4 to 6×10^5 seconds of computation, that is to say between 16 and 160 hours of computation. This estimate does not take into account the limitations of the memory: at each step, the storage in the computer memory of $\mathcal{O}(N^2)$ inner products as floating point numbers (4 bytes each) would require at least $(30000)^2 \times 4$ bytes, that is to say about 3.6 Gbytes. Without a super-computer, this implies using extensively the hard-drive for caching purposes, and this makes the computations much slower. One could indeed expect a couple of months of computations, which should be compared to the 2.5 seconds duration of the signal. On the other hand, the Fast Ridge Pursuit was run on a consumer PC running at 300 Mhz and equipped with 128 Mbytes of memory. It only took 200 seconds to get the result.

Figure 2 displays the decrease, in decibels, of the energy of the residual. It is faster with the Fast Ridge Pursuit than with the standard Gabor Matching Pursuit. This is not a trivial fact despite the chirp dictionary being more redundant than the Gabor dictionary. Actually, it is obvious that for a given sparseness (a number M of atoms), the chirp dictionary should give a better approximation quality *if we have at hand an algorithm to find the best M -atom approximation*. However the Pursuit strategy that we are following is sub-optimal, and there are examples [34] where choosing “better” atoms in a more redundant dictionary at each step yields worse approximations. It is thus important to observe that both Bultan's algorithm [14] and our Fast Ridge Pursuit with chirp atoms *do* provide a better approximation quality for a given sparseness than the Matching Pursuit with Gabor atoms. However the price to pay for this is the increased number of bits needed to describe the location (s_m, u_m, ξ_m, c_m) of the atoms. This is analogous to the situation where a codebook size of a vector quantizer is increased to allow better approximation: a clever encoding of the location of the vectors used in a given expansion is needed before using it for signal compression.

One can compare on Figure 3 the time-frequency distributions [15], [14] associated to the Gabor Matching Pursuit and Fast Ridge Pursuit decompositions of the signal. The display corresponds to a weighted linear combination

$$E_M[x](t, \omega) = \sum_{m=1}^M |\langle R^{m-1}x, g_m \rangle|^2 WV[g_m](t, \omega) \quad (23)$$

of the Wigner-Ville distribution of the atoms in the decom-




Fig. 2. Decay (in decibels) of the relative energy $\|R^m x\|^2 / \|x\|^2$ of the residual with the number m of iterations. Plain : Gabor Matching Pursuit. Bold : Fast Ridge Pursuit with chirp dictionary. One needs fewer chirp atoms than Gabor atoms to get the same approximation quality.

position

$$x(t) = \sum_{m=1}^M \langle R^{m-1} x, g_m \rangle g_m(t) + R^M(t). \quad (24)$$

It is focused on a time-frequency area wherein the vibrato occurs. The Gabor Matching Pursuit needs several constant-frequency atoms, located on the “path” of the instantaneous frequency, to decompose the *vibrato*. On the contrary, the Fast Ridge Pursuit decomposes it into only a few chirp atoms, whose instantaneous frequency is alternatively increasing and decreasing. Actually, both algorithms iterate 5000 times : at first, both algorithms select atoms that fit signal structures, and the energy of the residual decreases quite fast (see Figure 2); then, as the residual starts behaving like a random noise [22] with no emerging structure, the chosen atoms do no longer reflect signal structures, but simply decrease the energy of the residual as well as they can. What we observe is that the Gabor Matching Pursuit needs more atoms to represent signal structures than the Fast Ridge Pursuit.

VII. COMMENTS

We checked numerically that the fast estimate given by Proposition 1 fails for non-Gaussian windows (even for B-spline windows, which in some sense are close to Gaussian windows). Even if an analogous of Theorem 1 can be derived for such windows, the lack of analytic tools makes it difficult to derive an analogous of the fast and simple estimation procedure. It may be possible however to get fast estimates using *regression* [35] instead of linear interpolation to fit the local behavior of the spectrum around the best Gabor atom.

We do not cover in this article the theoretical question of the convergence of the Ridge Pursuit. One should notice that the convergence is in general not guaranteed by the fact that it is *stepwise more greedy* (the chosen chirp atom grabs more energy than the best Gabor atom) than the Gabor Matching Pursuit.

VIII. CONCLUSION

The Fast Ridge Pursuit algorithm iteratively decomposes an N -sample acoustic signal into M Gaussian chirp atoms




Fig. 3. Time-frequency distributions of a sound recording of size $N \approx 30000$ (total duration 2.5 seconds, sampling rate 11025 Hertz). Top : with $M = 5000$ iterations of Gabor Matching Pursuit. Bottom : with $M = 5000$ iterations of Fast Ridge Pursuit. The energy density is grey-coded relatively to its largest value, from $-45dB$ (white) to $0dB$ (black). The display is focused on a time-frequency region wherein the vibrato occurs visibly, while the whole time-frequency distribution would be for $0 \leq t \leq 2.5$ second and $0 \leq \frac{\omega}{2\pi} \leq 5500$ Hertz. Vertical lines (e.g. at time $t = 2.1$) correspond to short scale atoms that represent transients. Horizontal lines, associated to large scale constant frequency atoms, represent the resonance of the notes of the instruments of the orchestra. The vibrato is decomposed into several constant frequency atoms by the Gabor Matching Pursuit. On the contrary the Fast Ridge Pursuit decomposes it into only a few chirp atoms (see text).

with a computational cost $\mathcal{O}(MN)$. Thanks to its low computational complexity, the sparse structured representation of signals that it provides can become the basis for the implementation of a large variety of new processing tools.

Besides its potential use for signal compression, one of its most interesting features is its ability to decompose a signal into superimposed structures with different scale, frequency and chirp characteristics. Thanks to this decomposition property, it is possible to process separately the different parts (e.g. transients and steady parts) of a signal. Source separation can be achieved for sounds that have very different “chirp behavior”, such as a singer (with a strong vibrato) and an orchestra. Additionally, considering time-stretching or pitch-shifting applications, it is possible to keep the fine structure of transients while processing the harmonic part of a sound. Because they respect the structure of the transients, and as the chirp parameter enables them to fit more finely the phase of the signal, such pitch shifting schemes will generate less “pipe noise” than standard windowed Fourier transform based techniques. Moreover their implementation using the chirplet decomposition is straightforward.

ACKNOWLEDGMENTS

The author would like to thank E. Bacry and S. Mallat, from Ecole Polytechnique, for their encouragement and all

the interesting discussions they fed. The sound recording was kindly provided by Xavier Rodet, from IRCAM, whom I also wish to thank. All the numerical computations and figures were obtained using Last Wave [31], a freely available software under the GPL license.

APPENDIX

I. PROOF OF THE RIDGE THEOREM

In this appendix, we give a proof of Theorem 1. Building Taylor expansions of $b(t)$ and $\phi(t)$ near $t = u$, one can find $\theta_1(t), \theta_2(t) \in [u, u + t]$ such that $a(u + t) = a(u)e^{-b'(u)t - \frac{b''(u)}{2}t^2 - \frac{b'''(\theta_1(t))}{6}t^3}$ and $\phi(u + t) = \phi(u) + \phi'(u)t + \frac{\phi''(u)}{2}t^2 + \frac{\phi'''(\theta_2(t))}{6}t^3$. By changing variables, and using the definition of the Gaussian window $g(t) = \pi^{-1/4}e^{-t^2/2}$, we express $\langle x, g_{(s,u,\xi,c)} \rangle$ as

$$\frac{ae^{i\phi}}{(b''/\pi)^{1/4}} e^{\frac{(b')^2}{2b''}} \int_{-\infty}^{+\infty} g\left(\frac{1}{\sqrt{b''}}, -\frac{b'}{b''}, 0, 0\right)(t) \overline{g_{(s,0,\xi-\phi',c-\phi'')}}(t) e^{\frac{i^3}{6}(-b'''(\theta_1) + i\phi'''(\theta_2))} dt$$

Note that in all this proof we do not express the dependency of a, ϕ, ϕ', \dots on u , not that of θ_1, θ_2 on t . The integral can be rewritten as

$$\left\langle g_{(1/\sqrt{b''}, -b'/b'', 0, 0)}, g_{(s,0,\xi-\phi',c-\phi'')} \right\rangle + \epsilon(s, u, \xi, c) \quad (25)$$

where the error term $\epsilon(s, u, \xi, c)$ is

$$\int_{-\infty}^{+\infty} g\left(\frac{1}{\sqrt{b''}}, -\frac{b'}{b''}, 0, 0\right)(t) \overline{g_{(s,0,\xi-\phi',c-\phi'')}}(t) \left(e^{\frac{i^3}{6}(-b'''(\theta_1) + i\phi'''(\theta_2))} - 1 \right) dt.$$

Let us now bound the error term, using again the expression of the Gaussian window, and splitting the integral with a parameter η :

$$\begin{aligned} |\epsilon(s, u, \xi, c)| &\leq \frac{(b'')^{\frac{1}{4}}}{\pi^{\frac{1}{2}} s^{\frac{1}{2}}} \int_{-\infty}^{+\infty} e^{-\frac{t^2}{2s^2} - \frac{b''}{2}\left(t + \frac{b'}{b''}\right)^2} \\ &\quad \left| e^{\frac{i^3}{6}(-b'''(\theta_1) + i\phi'''(\theta_2))} - 1 \right| dt \\ &\leq \frac{(b'')^{\frac{1}{4}}}{\pi^{\frac{1}{2}} s^{\frac{1}{2}}} \left(\int_{|t| > \eta s} + \int_{|t| \leq \eta s} \right). \end{aligned}$$

The first part of the split integral is bounded by

$$\begin{aligned} &\int_{|t| > \eta s} e^{-\frac{t^2}{2s^2}} \left| e^{-\frac{(b')^2}{2b''}} e^{-b't - \frac{b''}{2}t^2 - \frac{i^3}{6}b'''(\theta_1)} e^{i\frac{t^3}{6}\phi'''(\theta_2)} - e^{-\frac{b''}{2}\left(t + \frac{b'}{b''}\right)^2} \right| dt \\ &= \int_{|t| > \eta s} e^{-\frac{t^2}{2s^2}} \left| \frac{a(u+t)}{a(u)} e^{-\frac{(b')^2}{2b''}} e^{i\frac{t^3}{6}\phi'''(\theta_2)} - e^{-\frac{b''}{2}\left(t + \frac{b'}{b''}\right)^2} \right| dt \\ &\leq \int_{|t| > \eta s} e^{-\frac{t^2}{2s^2}} \left(\frac{\|a\|_{\infty}}{|a(u)|} + 1 \right) dt \leq \frac{2\|a\|_{\infty}}{|a(u)|} \int_{|t| > \eta s} e^{-\frac{t^2}{2s^2}} dt \\ &\leq s \frac{2\|a\|_{\infty}}{|a(u)|} \frac{2}{\eta} e^{-\eta^2/2} \leq s \frac{\|a\|_{\infty}}{|a(u)|} \frac{4}{\eta} e^{-\eta^2/2}. \end{aligned}$$

Knowing that $\forall z \in \mathbb{C}$, $|e^z - 1| \leq e^{|z|} - 1 \leq |z|e^{|z|}$ we can bound the second part with

$$\begin{aligned} &\int_{|t| \leq \eta s} e^{-\frac{t^2}{2s^2}} \frac{|t|^3}{6} (|b'''(\theta_1)| + |\phi'''(\theta_2)|) e^{\frac{|t|^3}{6}(|b'''(\theta_1)| + |\phi'''(\theta_2)|)} dt \\ &\leq s \frac{\|b'''\|_{\infty} + \|\phi'''\|_{\infty}}{6} e^{\frac{\eta^3 s^3}{6}} (\|b'''\|_{\infty} + \|\phi'''\|_{\infty}) s^3 \sigma_3^3 \end{aligned}$$

We denote by $K = \|b'''\|_{\infty} + \|\phi'''\|_{\infty}$ and get from these two bounds that for all $\eta > 0$

$$|\epsilon(s, u, \xi, c)| \leq \frac{(b'')^{\frac{1}{4}} s^{\frac{1}{2}}}{\pi^{\frac{1}{2}}} \left(\frac{\|a\|_{\infty}}{a(u)} \frac{4e^{-\eta^2/2}}{\eta} + \frac{K s^3 \sigma_3^3}{6} e^{\frac{\eta^3 s^3}{6}} \right). \quad (26)$$

Choosing $\eta^3 = K^{-1} s^{-3}$ gives (11). To conclude the proof, we rewrite the first term of (25) as

$$e^{-i\phi' \frac{b'}{b''} + i\frac{\phi''}{2} \left(\frac{b'}{b''}\right)^2} \left\langle g_{\left(\frac{1}{\sqrt{b''}}, u - \frac{b'}{b''}, \phi' - \phi'' \frac{b'}{b''}, \phi''\right)}, g_{(s,u,\xi,c)} \right\rangle.$$

II. REAL-VALUED ATOMS

Let $\tilde{g}_{\gamma} = (g_{\gamma} - \overline{g_{\gamma}} \overline{g_{\gamma}}) / \left(1 - |\overline{g_{\gamma}} g_{\gamma}|\right)^2$ and $\mathcal{V}_{\gamma} = \text{Span}(g_{\gamma}, \overline{g_{\gamma}})$. One can check that $\{\tilde{g}_{\gamma}, \overline{\tilde{g}_{\gamma}}\}$ is the dual basis of $\{g_{\gamma}, \overline{g_{\gamma}}\}$. Thus for all $x \in \mathcal{H}$

$$\begin{aligned} P_{\mathcal{V}_{\gamma}} x &= \langle x, \tilde{g}_{\gamma} \rangle g_{\gamma} + \langle x, \overline{\tilde{g}_{\gamma}} \rangle \overline{g_{\gamma}} \\ \|P_{\mathcal{V}_{\gamma}} x\|^2 &= \frac{2 \left(|\langle x, g_{\gamma} \rangle|^2 - \Re(\langle g_{\gamma}, \overline{g_{\gamma}} \rangle \langle x, g_{\gamma} \rangle) \right)}{1 - |\overline{g_{\gamma}} g_{\gamma}|^2} \end{aligned}$$

where $\Re(z)$ denotes the real part of $z \in \mathbb{C}$. For x real-valued and $\gamma = (s, u, \xi, c)$, the first equality can be rewritten $P_{\mathcal{V}_{(s,u,\xi,c)}} x = \langle x, g_{(s,u,\xi,c,\phi)} \rangle g_{(s,u,\xi,c,\phi)}$ with $\langle x, g_{(s,u,\xi,c,\phi)} \rangle = \|P_{\mathcal{V}_{\gamma}} x\|$ and $e^{i\phi} = \frac{\langle x, g_{\gamma} \rangle}{|\langle x, g_{\gamma} \rangle|}$. The value $\langle \overline{g_{(s,u,\xi,c)}} , g_{(s,u,\xi,c)} \rangle$ can be computed up to an arbitrary precision with a cost $\mathcal{O}(1)$, thanks to an analytic expression [14] [24]. Once $\langle x, g_{\gamma} \rangle$ is known, so is its complex conjugate $\langle x, \overline{g_{\gamma}} \rangle$, thus $\phi(s, u, \xi, c)$ and $\langle x, g_{(s,u,\xi,c,\phi)} \rangle$ can be computed in $\mathcal{O}(1)$.

REFERENCES

- [1] I. Daubechies, "Orthonormal bases of compactly supported wavelets," *Commun. on Pure and Appl. Math.*, vol. 41, pp. 909–996, Nov. 1988.
- [2] S. Mallat, "A theory for multiresolution signal decomposition : the wavelet representation," *IEEE Trans. Pattern Analysis and Machine Intelligence*, pp. 674–693, July 1989.
- [3] G. Beylkin, R. Coifman, and V. Rokhlin, "Fast wavelet transforms and numerical algorithms," *Commun. on Pure and Appl. Math.*, vol. 44, pp. 141–183, 1991.
- [4] S. Jaffard, "Pointwise smoothness, two microlocalization and wavelet coefficients," *Publicacions Matemàtiques*, vol. 35, pp. 155–168, 1991.
- [5] S. Mallat and W.L. Hwang, "Singularity detection and processing with wavelets," *IEEE Trans. Inform. Theory*, vol. 38, no. 2, pp. 617–643, Mar. 1992.
- [6] S. Mallat and S. Zhong, "Characterization of signals from multiscala edges," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 40, pp. 2464–2482, July 1992.
- [7] R. R. Coifman and Y. Meyer, "Remarques sur l'analyse de Fourier à fenêtre," *Comptes-Rendus Acad. Sci. Paris(A)*, vol. 312, pp. 259–261, 1991.
- [8] R. Coifman and M.V. Wickerhauser, "Entropy-based algorithms for best basis selection," *IEEE Trans. Inform. Theory*, vol. 38, no. 2, pp. 713–718, Mar. 1992.

- [9] B. Torrèsani, "Wavelets associated with representations of the affine Weyl-Heisenberg group," *J. Math. Phys.*, vol. 32, pp. 1273–1279, May 1991.
- [10] S. Qian and D. Chen, "Signal representation using adaptive normalized Gaussian functions," *Signal Process.*, vol. 36, no. 1, pp. 1–11, 1994.
- [11] H. K. Kwok and D. L. Jones, "Improved FM demodulation in a fading environment," in *Proc. IEEE Conf. Time-Freq. and Time-Scale Anal. (TFTS'96)*, Paris, June 1996, pp. 9–12.
- [12] R. G. Baraniuk and D. L. Jones, "Shear madness : New orthonormal bases and frames using chirp functions," *IEEE Trans. Signal Process. Special Issue on Wavelets in Signal Processing*, vol. 41, no. 12, pp. 3543–3548, Dec. 1993.
- [13] S. Mann and S. Haykin, "The chirplet transform : Physical considerations," *IEEE Trans. Signal Process.*, vol. 43, no. 11, pp. 2745–2761, Nov. 1995.
- [14] A. Bultan, "A four-parameter atomic decomposition of chirplets," *IEEE Trans. Signal Process.*, vol. 47, no. 3, pp. 731–745, Mar. 1999.
- [15] S. Mallat and Z. Zhang, "Matching pursuit with time-frequency dictionaries," *IEEE Trans. Signal Process.*, vol. 41, no. 12, pp. 3397–3415, Dec. 1993.
- [16] P. Flandrin, *Temps-Fréquence*, Hermes, Paris, France, 1993.
- [17] S. Mallat, *A Wavelet Tour of Signal Processing*, Academic Press, 1998.
- [18] G.H. Watson and K. Gilholm, "Signal and image feature extraction from local maxima of generalized correlation," *Pattern Recognition*, vol. 31, no. 11, pp. 1733–1745, 1998.
- [19] I. Daubechies, *Ten Lectures on Wavelets*, SIAM, 1992.
- [20] P. J. Huber, "Projection pursuit," *The Annals of Statistics*, vol. 13, no. 2, pp. 435–475, 1985.
- [21] L.K. Jones, "On a conjecture of Huber concerning the convergence of PP-regression," *The Annals of Statistics*, vol. 15, pp. 880–882, 1987.
- [22] G. Davis, S. Mallat, and M. Avellaneda, "Adaptive greedy approximations," *Constr. Approx.*, vol. 13, no. 1, pp. 57–98, 1997.
- [23] Y.C. Pati, R. Rezaifar, and P.S. Krishnaprasad, "Orthonormal matching pursuit : recursive function approximation with applications to wavelet decomposition," in *Proceedings of the 27th Annual Asilomar Conf. on Signals, Systems and Computers*, Nov. 1993.
- [24] A. Papoulis, *The Fourier integral and its applications*, McGraw-Hill classic textbook reissue series. McGraw-Hill Publisher Co., 1987.
- [25] R. Gribonval, *Approximations non-linéaires pour l'analyse de signaux sonores*, Ph.D. thesis, Université Paris IX Dauphine, Sept. 1999.
- [26] N. Delprat, B. Escudé, P. Guillemain, R. Kronland-Martinet, Ph. Tchamitchian, and B. Torrèsani, "Asymptotic wavelet and Gabor analysis : Extraction of instantaneous frequency," *IEEE Trans. Inform. Theory*, vol. 38, no. 2, pp. 644–664, Mar. 1992.
- [27] J. S. Marques and L. B. Almeida, "A background for sinusoid based representation of voiced speech," in *Proc. Int. Conf. Acoust. Speech Signal Process. (ICASSP'86)*, Tokyo, 1986, pp. 1233–1236.
- [28] J. S. Marques and L. B. Almeida, "Frequency-varying sinusoidal modeling of speech," *IEEE Trans. Speech and Audio Process.*, vol. 37, no. 5, pp. 763–765, May 1989.
- [29] F. Bergeaud, *Représentations adaptatives d'images numériques*, Matching Pursuit, Ph.D. thesis, Ecole Centrale Paris, 1995.
- [30] F. Bergeaud and S. Mallat, "Matching pursuit : Adaptive representations of images and sounds," *Computational and Applied Mathematics*, vol. 15, no. 2, Oct. 1996.
- [31] E. Bacry, *LastWave software*, <http://wave.emap.polytechnique.fr/soft/LastWave/>, Freely available (GPL license).
- [32] M.-A. Dalbavie, "“Marc-André Dalbavie”," in *Compositeurs d'Aujourd'hui*. IRCAM, 1999:1–93.
- [33] X. Rodet, "Time-domain formant-wave functions synthesis," in *Spoken Language Generation and Understanding*, J.C. Simon, Ed., C: Mathematical and Physical Sciences, chapter 4-Speech Synthesis, pp. 429–441. D. Reidel Publishing Company, 1980.
- [34] R. Gribonval, "A counter-example to the general convergence of partially greedy algorithms," *to appear in J. Approx. Theory*, 2001.
- [35] C.M. McIntyre and D.A. Dermott, "A new fine-frequency estimation algorithm based on parabolic regression," in *Proc. Int. Conf. Acoust. Speech Signal Process. (ICASSP'92)*, 1992, pp. 541–544.

R. Gribonval graduated from École Normale Supérieure, Paris, France, in 1997. He received the Ph.D. degree in applied mathematics from the Université Paris-IX Dauphine, Paris, France, in 1999. In 2000 he was a visiting scholar at the Industrial Mathematics Institute (IMI) in the Department of Mathematics, University of South Carolina, SC. He is currently a Research Associate with the French National Center for Computer Science and Control (INRIA) at IRISA, Rennes, France. His current research interests are in adaptive techniques for the representation and classification of audio signals with redundant systems.