Mon. Not. R. Astron. Soc. **308**, 876–886 (1999)

Band-limited interpolation applied to the time series of rapidly oscillating stars

S. Roques,★ B. Serre and N. Dolez
Laboratoire d'Astrophysique, Observatoire Midi-Pyrénées, 14 Avenue Edouard Belin, F-31400 Toulouse, France

Accepted 1999 April 29. Received 1999 March 5; in original form 1997 November 14

ABSTRACT

We reconstruct the temporal Fourier spectra obtained from incomplete records of the luminosities of rapidly oscillating stars. This reconstruction involves a deterministic procedure related to 'band-limited interpolation'. We show that the stability of the underlying deconvolution process is governed by the choice of the frequency intervals over which the restoration has to be achieved. A 'time–frequency analysis' provides a progressive constructive process allowing us to choose these intervals, to measure the frequencies and the fine structure of the driven modes of the star, and to control the error propagation. The method, which is the first step towards multisite deconvolution, is described in detail, and numerical examples for the DAV white dwarf GD 154 are presented to illustrate the application of the algorithm. The restored frequencies are compared with data obtained during a Whole Earth Telescope campaign.

Key words: methods: data analysis – techniques: image processing – stars: individual: GD 154 – stars: oscillations – stars: variables: other – white dwarfs.

1 INTRODUCTION

Many classes of variable stars are known that have oscillation periods short enough to allow astronomers to record a large number of cycles in a single night. They include, for example, three classes of white dwarf (DAV, DBV and PG 1159 stars), nuclei of planetary nebulae, RoAp stars, short-period δ Scuti stars, and the newly discovered class of EC 14026 stars (pulsating hot subdwarfs: see Charpinet et al. 1996; Kilkenny et al. 1997). These objects are interesting targets from the asteroseismological point of view, as they are multiperiodic, and the observation of rapid oscillations allows an accurate determination of the frequencies.

Using classical photomultipliers, and (more and more frequently) CCD cameras, fast photometry is a well-suited observational technique for this kind of star. It yields equally (or almost equally) spaced time series representing the luminosity variations of the variables.

The frequencies of the driven modes represent an important source of information: first, they provide tests for determining whether or not the time series was generated by a set of linear eigenmodes of the star; secondly, the fine structure of each mode yields the rotation period of the star; thirdly, their frequency spacings provide its mass and the composition of its layers (Koester & Chanmugan 1990). Usually, to gain information accurate enough for asteroseismology, in terms of signal-to-noise ratio and frequency resolution, requires as long an observing run as possible, and single-night observations are insufficient; therefore rapidly oscillating stars are more and more often observed via worldwide networks (Whole Earth Telescope: e.g. Nather et al. 1990; STEllar PHotometry International: e.g. Michel et al. 1992; Delta Scuti Network: e.g. Breger, Handler & Nather 1995).

Unfortunately, during the recording of the luminosity variations of the star, diurnal cycles or uncooperative weather can cause incomplete coverage of the records; the length and the irregularity of the gaps yield degradations of the Fourier spectrum of the data, which reduce the expected frequency resolution. As it is practically impossible to eliminate data gaps from the observations themselves, one is then forced to minimize their effects on the Fourier spectra by designing better signal restoration algorithms, allowing us to reach the theoretical resolution. These aspects are presented in Section 2.

Sections 3 and 4 are devoted to the reconstruction of the temporal Fourier spectra, from a deterministic procedure related to 'band-limited interpolation' and to 'time–frequency analysis'. Band-limited interpolation has already been applied to the resolution of two-dimensional deconvolution problems, and in particular for restoration of *Hubble Space Telescope* (*HST*) images (Bouyoucef, Fraix-Burnet & Roques 1997). The originality of the method is in taking into account the presence of noise in the input data and providing estimations of the quadratic error of the reconstruction.

We show that, in such a problem, the stability of the underlying

★ E-mail:roques@obs-mip.fr

deconvolution process is in particular governed by the choice of the frequency intervals over which the restoration has to be achieved. At the same time, the necessity of getting information about the energy together with the lifetime of a given feature in the Fourier spectrum leads us to time–frequency analysis, used in parallel with Fourier analysis. In this context, the 'matching pursuit' algorithm (Mallat & Zhang 1993) plays a crucial role. The associated constructive process allows us to detect and to characterize the time–frequency components of the light curve one by one, from the highest energy one to the lowest. Then we can reconstruct the spectrum of the star step by step, under the condition that it can be broken up into frequency domains, whereas we show that a global reconstruction would have been unstable at the same resolution.

Finally, in Section 5, we apply the method to the DAV white dwarf GD 154, and discuss its degree of confidence obtained from an error analysis, and from a comparison with data obtained during a Whole Earth Telescope (WET) campaign.

2 THE NATURE OF THE DATA

To take into account the maximum reliable frequency information contained in the observation of a given rapidly oscillating star, it is necessary to have as long a light curve as possible, since the resolution is inversely proportional to the observation length. Under this constraint, it is often the case that a single-night observation does not display sufficient resolution to extract directly, from the brightness record of the star, information of astrophysical interest; however, in that case of a limited time span, there are few artefacts in the Fourier spectrum and the signal-to-noise ratio is large enough.

One way to circumvent the problem of diurnal breaks is to put the star in a coordinated multi-site observing campaign between telescopes distributed over several longitudes, and afterwards to reduce records to individual normalized light curves, as in for example the WET campaigns for white dwarfs (Nather et al. 1990). In practice, owing to weather conditions or equipment malfunctions, the coverage of the time series that can be obtained usually remains incomplete (80 per cent maximum). As a result, each feature in the spectrum is accompanied by sidelobes and satellite peaks, the nature of which closely depends on the distribution of observing intervals. As they can be of substantial amplitude, they can lead to confusion of features arising from the oscillation modes with those resulting from the segmented nature of the observing window. Thus the multi-site observation, more than the one-night one, may generate systematic errors (e.g. sidelobes) in the temporal spectra, although the resolution is improved.

In this paper, we deal with the two types of data that are equivalent from a methodological point of view. In both cases, we have experimental light curves of the form

$$\psi(t) = \phi(t)w(t) + \text{noise}, \tag{1}$$

where $w(t)$ is the 'observing window function' characteristic of the observing intervals, $\psi(t)$ is the experimental light curve at some level of resolution, and $\phi(t)$ is the intensity of the source at the same level of resolution but without gaps. This resolution level is imposed by the sampling of the record. The dual data of $\psi(t)$,

$$\hat{\psi}(u) = (\hat{\phi} \star \hat{w})(u) + \hat{\xi}(u), \tag{2}$$

in Fourier space are then the result of a convolution operation; the related kernel of convolution $\hat{w}(u)$ is the 'impulse response', Fourier transform of the observing window function. The additive term $\hat{\xi}(u)$ includes random and systematic errors, for instance errors on sampling, or noise arising from detectors, atmosphere, etc. (see Section 3 and Appendix A).

We are facing a deconvolution problem in Fourier space (equation 2), where we try to reconstruct, at least to some extent, the spectrum that would have been obtained with a non-discontinuous (the case of a multi-site observation) or wider (the case of a one-night observation) observing window function.

The problem is then equivalent to interpolating or extrapolating the time series in the gaps. Unfortunately, however, in these terms, the problem is ill-conditioned, and the solutions are unstable with respect to small variations of the initial data (Tikhonov & Arsenin 1977).

To circumvent this instability, one has to consider that the Fourier data $\hat{\psi}$ are (practically) of bounded support in some frequency domain $D_{\rm f}$.

The stabilization procedure consists then of performing interpolation (in the gaps) and some *partial* extrapolation [equivalent to interpolation by (i) padding the beginning and the end of the time series with zeroes; and (ii) interpolating as if it was a gap]. In other words, we have to interpolate, in some temporal region, the inverse Fourier transform of a function having its support in some given frequency region $D_{\rm f}$. This is referred to as 'interpolation of band-limited functions' (Youla 1978; Sanz & Huang 1983). Two domains play an essential part in the stability of the procedure: the size of $D_{\rm f}$ (support of the Fourier spectrum); and the size of the interpolation time domain together with the distribution of data gaps.

3 REGULARIZATION OF THE INVERSE PROBLEM

3.1 The spectrum to be reconstructed

The data used in this paper have been obtained by a two-channel photometer, recording the pulsating star and a comparison star. After acquisition, the raw data are reduced according to a standard procedure to remove the effects of background sky brightness, atmospheric extinction and motion of the Earth around the barycentre of the Solar system (Clemens 1993). As the quality of the run was not extremely good, owing to atmospheric fluctuations affecting the sky transparency, it has been necessary to use the comparison stars to correct those effects: this procedure efficiently removes some systematic errors, but unfortunately increases the level of noise in the affected parts of the time series (see Fig. 1).

Owing to these noise sources and to the different types of error, it is preferable to give up trying to restore the original spectrum $\hat{\phi}(u)$ at its highest level of resolution. This leads us to define $\hat{\phi}_s(u)$, the 'spectrum to be reconstructed', as a smoothed version of the ideal spectrum $\hat{\phi}(u)$ by a relation of the form

$$\hat{\phi}_s(u) = (\hat{\phi} \star \hat{s})(u), \tag{3}$$

where the support of $\hat{\phi}_s$ is the frequency domain $D_{\rm f}$.

The smoothing function $\hat{s}(u)$ can be seen as a synthetic impulse response, Fourier transform of a synthetic window function. It is designed so that its support is as small as possible and its inverse Fourier transform is small, in the mean-squares sense, outside a synthetic aperture $W_{\rm r}$ regularizing the support W of the observing

window function, and in particular wider than this function (see Fig. 2). The existence of such a function, as well as the possibility of obtaining accurate approximations to its values, does not raise any particular difficulties (Slepian 1964).

The choice of this synthetic window function is of fundamental importance because it defines the resolution limit of the reconstruction process. The size of its support varies as the inverse of the size of W_r. Intuitively, the smaller the support of $\hat{s}(u)$ with respect to $\hat{w}(u)$ (or the greater W_r with respect to W), the greater the gain in frequency resolution. Unfortunately, in that case, numerical instabilities owing to the ill-conditioned nature of the matrix of the deconvolution system appear, and the solution may be unstable in the sense that a small variation in the data can lead to a very large difference in the solutions.

3.2 Regularization function

The ratio of the signal to the noise defines the point-to-point signal-to-noise ratio in direct space:

$$S/N(t) = \frac{|\psi(t)|}{n(t)}, \tag{4}$$

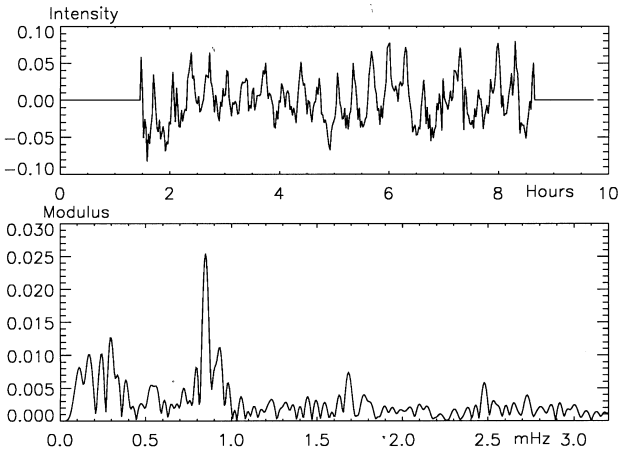

Figure 1. Top: an example of a normalized centred light curve, after the reduction procedure as described in Section 3.1. Some localized defects cannot be completely removed and increase the overall level of noise in the spectral domain (here GD 154 observed on 1991 May 17 at Mauna Kea Observatory). The length of the observing window is 7 h. Bottom: the Fourier transform (in modulus) of the signal. The amplitude of modes around 1 mHz is well above the noise level, whereas it is more difficult to extract information in the range 2–3 mHz.

where $n(t)$ characterizes the error–object bound such that

$$|\psi(t) - \phi(t)w(t)| \leq n(t). \tag{5}$$

In practice, the determination of $n(t)$ should of course result from an appropriate (theoretical or experimental) estimation of the various types of errors. Since this analysis may be very complex, it can be simplified by using the statistical model of soft-thresholding proposed by Donoho & Johnstone (1994) and summarized in Appendix A.

Depending on the values of the signal-to-noise ratio, the principal features of ϕ_s are in practice roughly given by

$$\tilde{\psi}(t) = \tilde{s}(t)\psi(t), \tag{6}$$

where $\tilde{s}(t) = s(t)/w(t)$ when $w(t) \neq 0$ and $S/N(t)$ is sufficiently large (for instance, greater than some threshold $\alpha \simeq 1$), and 0 otherwise (this means that the corresponding signal information is discarded a priori). This operation is a kind of adaptive filter: first a coarse deconvolution in frequency is carried out when dividing the data by $w(t)$, and the frequency result is convolved by the synthetic impulse response $\hat{s}(u)$ which smoothes the restored spectrum [this is done with the multiplication by $s(t)$ in time space]. Note that, if one chooses as $\hat{s}(u)$ a box-car function, the solution obtained is the so-called Bracewell principal solution (Bracewell 1955). In the domain where the signal-to-noise ratio is too bad, we give up processing information. This approximation is good as a first step, but can be improved depending on the nature of the data. We can now define the reconstructed spectrum as the function that minimizes in the mean-squares sense a functional of the form (Lannes, Roques & Casanove 1987c)

$$q(\phi) = \|g(t)(\tilde{\psi}(t) - \phi(t))\|^2. \tag{7}$$

The method of 'conjugate gradients' (Hestenes & Stiefel 1952) is particularly well suited to solving this equation.

The 'regularization function' $g(t)$ is a weight function characteristic of the amount of weighted interpolation to be done in direct space, and is defined in relation to $S/N(t)$ and to W_r. In particular, it satisfies the following conditions.

(i) $g(t) = 1$ outside W_r. This condition defines the regularization principle under consideration: we give up restoring the spectrum $\hat{\phi}(u)$ at its highest level of resolution, in other words we give up synthesizing the whole time series $\phi(t)$ (total extrapolation is forbidden). Practically, the condition $g(t) = 1$ means that the spectrum $\hat{\phi}(u)$ is known, since it is set equal to zero.

(ii) $g(t) = 0$ in the parts of W_r where $S/N(t) \leq \alpha$. This

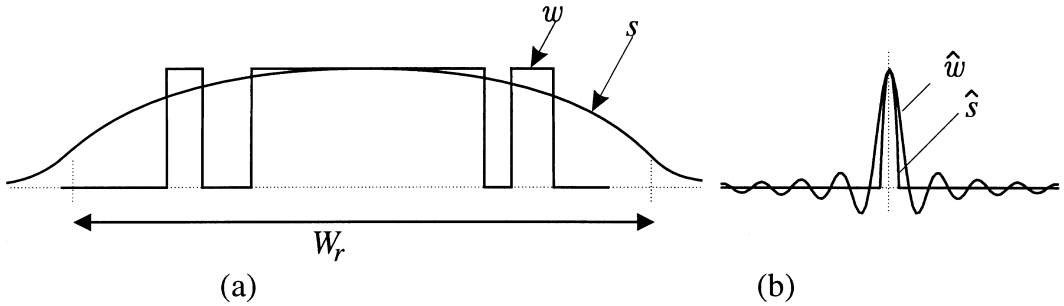

Figure 2. Resolution parameters in direct (a) and in Fourier (b) spaces. The support of $\hat{s}(u)$ is as small as possible, the idea being to have the best possible resolution. The energy of its inverse transform $s(t)$ is concentrated in the interval W_r; to this end, we impose a condition of the form $\|w_r s\|^2/\|s\|^2 = \chi^2$, where w_r is the characteristic function of W_r and χ^2 the energy rate of $s(t)$ in W_r (say, $\chi = 0.95$). The function $\hat{s}(u)$ is normalized to 1.

condition corresponds in particular to the total interpolation of the time gaps in W_r.

(iii) $g(t)$ is an increasing function of the signal-to-noise ratio, if $S/N(t) \geq \alpha$. This allows us to take into account, for instance, the fact that some parts of the signal are of higher quality than others, which is the case in particular for WET campaigns when we deal with different sites and different telescopes, or when uncooperative weather occurs during a part of the night.

For example, the regularization function may be defined as follows:

$$g(t) = \begin{cases} 1 & \text{outside } W_r, \\ 1 & \text{in } W_r \text{ if } S/N(t) \geq \beta \\ & \text{(information entirely reliable),} \\ \dfrac{S/N(t) - \alpha}{\beta - \alpha} & \text{in } W_r \text{ if } \alpha \leq S/N(t) < \beta \\ & \text{(information partially reliable),} \\ 0 & \text{in } W_r \text{ if } S/N(t) < \alpha \\ & \text{(information not reliable).} \end{cases} \quad (8)$$

Here, β is the threshold over which the information is considered as entirely reliable (typically we can choose $\beta = 3$ or 4; this means that, to keep all the information, the signal has to be at least three or four times more significant than the noise).

In such a problem, the least-squares solution of the deconvolution problem (equation 7) is solution of the 'normal equation':

$$A^*A\hat{\phi}(u) = A^*\hat{g}(u)\hat{\tilde{\psi}}(u), \quad (9)$$

where the so-called 'imaging operator' is $A^*A = vFg^2F^*$. Here, F and F^* stand for direct and inverse Fourier transform, and v is the characteristic function of the deconvolution frequency support D_f.

The smallest eigenvalue μ of A^*A conditions the stability of the reconstruction problem. It can be analytically estimated as a fuction of v and g. Indeed, this eigenvalue is a function of the 'interpolation parameter':

$$\eta = \left[\int_{D_f} v(x)\,\mathrm{d}x\right]^{1/2} \left\{\int_{W_r} [1 - g^2(u)]\,\mathrm{d}u\right\}^{1/2}, \quad (10)$$

characterizing the amount of weighted interpolation to be performed both in real and in Fourier space. This term is the geometric mean of the measures of D_f and W_r, and is thus a size parameter without dimension. Provided that η is sufficiently small, the minimum eigenvalue of A^*A can be obtained from a series of the form

$$\mu(\eta) = 1 - \eta^2 \sum_{m=0}^{\infty} (-1)^m \varrho_m \eta^{2m} \quad (\varrho_0 = 1), \quad (11)$$

where the ϱ_ms depend on v and g and are of the 'moment of inertia' kind relative to W_r (Lannes, Roques & Casanove 1987a).

The minimum eigenvalue μ of A^*A occurs in the expression of an upper bound Θ of the relative reconstruction error:

$$\frac{\|\Delta\phi\|}{\|\phi\|} \leq \Theta, \quad \text{where } \Theta = \frac{1}{\sqrt{\mu}} \frac{\|g\tilde{s}\mathfrak{n}\|}{\|\tilde{\psi}\|}. \quad (12)$$

In most cases, the error on the signal is an increasing function of the size of W_r and D_f. Thus the smaller are these intervals, the more reliable is the reconstruction. Then, by suitably choosing W_r and D_f, the size of the error can be acceptably small.

This choice, based on the estimation of Θ and executable in an interactive manner before reconstruction, allows us to realize the best possible compromise between resolution and reliability. So, since a given situation can be recognized a priori as being more or less well-conditioned, we can have an idea of the stability of the problem before its implementation. From this point of view, the indication provided by the moments of inertia ϱ_m is also very useful. For example, for a given value of η, the robustness of the reconstruction process is increased whenever the time gaps are well distributed over W_r.

If the error happens to be too large, the ill-conditioned problem is restated either at a lower level of resolution by reducing W_r, or, better, by improving the choice of the frequency support of deconvolution D_f, a choice which may be non-trivial. Nevertheless, this last point is the direction in which we want to go to stabilize the problem, because it does not imply a decrease of the resolution.

4 TIME–FREQUENCY ANALYSIS

When we try to perform deconvolution on the whole frequency support, the problem is generally ill-conditioned, or else it must be restated at a lower level of resolution. We are then led to choose particular intervals over which the deconvolution will be achieved. The idea is to build the reconstruction space step by step, by using a progressive constructive process allowing us to isolate these supports. However, it turns out that the choice of these intervals may be difficult. Indeed, if the choice of a support for low-frequency peaks is rather easy because they have large amplitudes, it is more difficult for high frequencies where the amplitudes of the peaks decrease more quickly than the noise. As a result, the intervals including low-amplitude peaks are hard to select by eye (Fig. 1).

Moreover, we search for oscillations arising almost everywhere in the signal, because they are characteristic of the structural properties of the star. So one understands the necessity of getting information about the lifetime of a given peak of the Fourier spectrum, the aim being to give us the possibility of eliminating the oscillations with too short a lifetime.

Then the particular nature of this problem leads us naturally to time–frequency analysis as a way of choosing these supports.

4.1 Matching pursuit algorithm

In the present stage of this work, the matching pursuit algorithm, introduced by Mallat & Zhang (1993), plays a decisive role. This algorithm allows us to choose, in a given redundant finite dictionary of time–frequency waveforms, a set of vectors that match the signal as well as possible.

The dictionary $\mathscr{D}$ is defined as a family of time–frequency functions obtained by dilating, modulating and translating a single real even function $k(t) \in L^2(\boldsymbol{R})$. The 'atoms' (elements) of the dictionary are defined by

$$k_\nu(t) = \frac{1}{\sqrt{a}} k\left(\frac{t-b}{a}\right) \mathrm{e}^{\mathrm{i}\omega t}, \quad (13)$$

where a is the dilation scale, b the translation parameter and ω the frequency modulation. One defines $\nu = (a, b, \omega)$ as the atom index

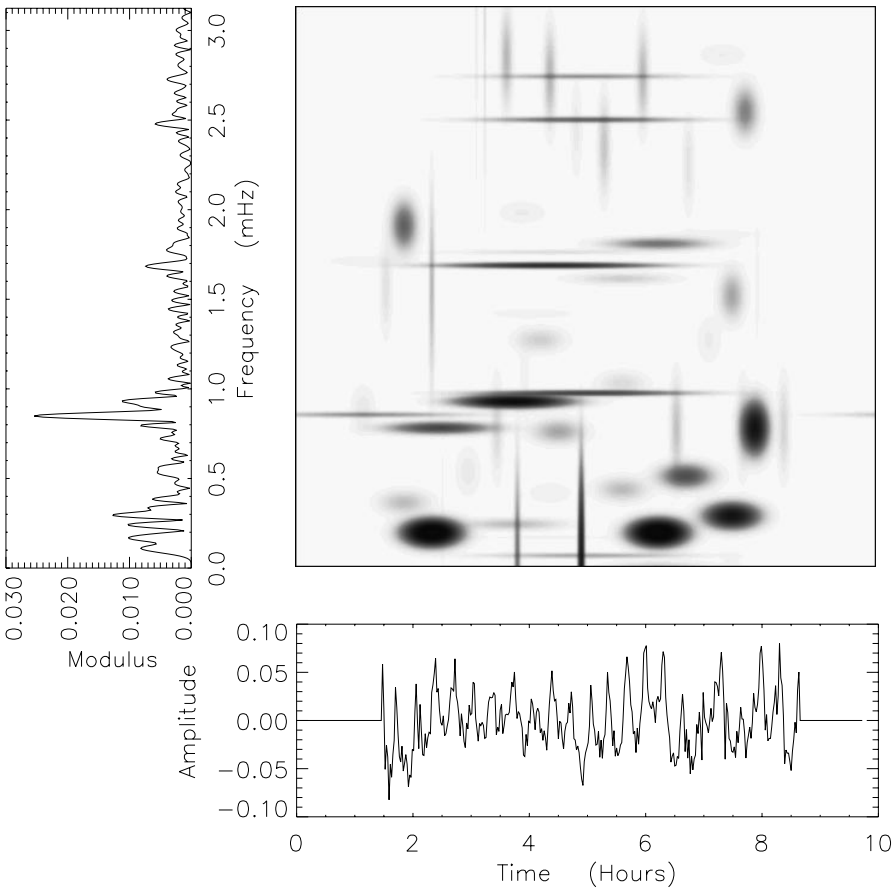

Figure 3. The time–frequency energy distribution obtained by using the matching pursuit algorithm on the signal presented in the bottom panel. The x-axis is time and the y-axis corresponds to the frequency. The Fourier transform is recalled on the left. A Gabor function dictionary is used, and the algorithm converges after 450 iterations.

in the dictionary. Note that the factor $1/\sqrt{a}$ normalizes to 1 the $L^2(\boldsymbol{R})$ norm of $k_\nu(t)$. If the window $k(t)$ is Gaussian, the joint time–frequency localization of all the atoms is a minimum, and in that case $k_\nu(t)$ is called a 'Gabor function'.

A matching pursuit algorithm computes adaptive signal representations: it expands any signal in a set of atoms selected among the extremely redundant dictionary $\mathscr{D}$, to match its components as well as possible, through iterated one-dimensional projections.

In our particular case, it is the light curve $\psi(t)$, for which the mean has first been subtracted so that it is a fluctuation about zero, that is approximated with a single vector k_{ν_0} chosen in the dictionary $\mathscr{D}$ such that $|\langle\psi(t), k_{\nu_0}(t)\rangle|$ is as large as possible. The light curve is then decomposed into this form:

$$\psi(t) = \langle\psi(t), k_{\nu_0}(t)\rangle k_{\nu_0}(t) + R\psi(t), \tag{14}$$

where $R\psi(t)$ is the residual vector after approximating $\psi(t)$ in the 'direction' $k_{\nu_0}(t)$. Clearly, $k_{\nu_0}(t)$ is orthogonal to $R\psi(t)$, and hence one has

$$\|\psi(t)\|^2 = |\langle\psi(t), k_{\nu_0}(t)\rangle|^2 + \|R\psi(t)\|^2. \tag{15}$$

The main idea of the matching pursuit is to sub-decompose the residue $R\psi(t)$, by finding a vector $k_{\nu_1}(t)$ that matches it as well as possible, as was done for $\psi(t)$. Each time, the procedure is repeated on the residue that is obtained:

$$\psi(t) = \langle\psi(t), k_{\nu_0}(t)\rangle k_{\nu_0}(t) + R\psi(t), \tag{16}$$

$$\mathrm{R}\psi(t) = \langle R\psi(t), k_{\nu_1}(t)\rangle k_{\nu_1}(t) + R^2\psi(t),$$

$$\vdots$$

$$R^n\psi(t) = \langle R^n\psi(t), k_{\nu_n}(t)\rangle k_{\nu_n}(t) + R^{n+1}\psi(t).$$

It is easy to determine a convergence criterion of the algorithm, by examining the decrease of the norm of the residue. Finally, the signal is decomposed into

$$\psi(t) = \sum_{i=0}^{\infty} \langle R^i\psi(t), k_{\nu_i}(t)\rangle k_{\nu_i}(t), \tag{17}$$

where the atoms $k_{\nu_i}(t)$ are the ones that match the signal structures as well as possible. We can then build a hierarchy of coherent structures $(k_{\nu_0}(t),\ k_{\nu_1}(t),\ \ldots,\ k_{\nu_n}(t))$, yielding a time–frequency energy distribution of the signal.

An energy conservation theorem results from equation (15):

$$\|\psi(t)\|^2 = \sum_{i=0}^{\infty} |\langle R^i\psi(t), k_{\nu_i}(t)\rangle|^2. \tag{18}$$

The energy density of $\psi(t)$ in the time–frequency plane (t, u) is

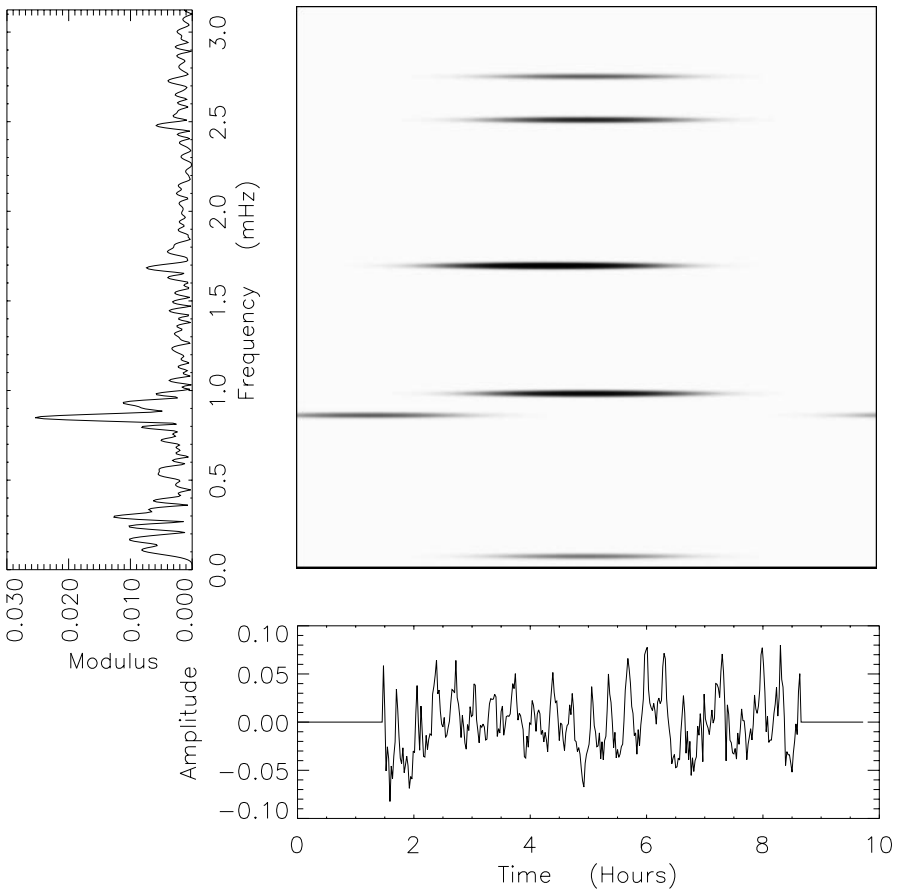

Figure 4. The time–frequency energy distribution of the first five coherent structures of the signal. These five most important groups allow us to choose the frequency intervals defining the deconvolution supports.

defined by

$$(E\psi)(t,u) = \sum_{i=0}^{\infty} |\langle R^i\psi(t), k_{\nu_i}(t)\rangle|^2 (Wk_{\nu_i})(t,u), \tag{19}$$

where $(Wk_{\nu_i})(t,u)$ is the Wigner–Ville distribution (Ville 1948) defined as follows:

$$(Wk_\nu)(t,u) = \int_{-\infty}^{\infty} k_\nu\left(t+\frac{\tau}{2}\right) k_\nu^\star\left(t-\frac{\tau}{2}\right) \mathrm{e}^{-\mathrm{i}u\tau}\,\mathrm{d}\tau. \tag{20}$$

In practice, it is this energy density $E\psi$ (equation 19) that is represented in the time–frequency diagrams (see e.g. Fig. 3). Note that this density does not include the interference terms of the Wigner–Ville distributions, because it is computed from the atomic decomposition of $\psi(t)$.

In contrast to decompositions on to orthonormal bases which are barely interpretable, or to the algorithm of Wickerhauser & Coifman (1992) which selects, in a global way, the basis that is best adapted to the signal properties, the matching pursuit decomposition is a constructive process which allows us to detect and characterize the time-frequency components one by one, from the highest energy one to the lowest.

4.2 Cooperation deconvolution: matching pursuit

In a matching pursuit diagram, we can choose to select only some atoms representing the most coherent structures of the signal; these atoms will appear as long (in time) pictures in the time–frequency plane (horizontal atoms). On the other hand, as the noise of the time series will not correlate well with any particular dictionary element, its information will be diluted and then sub-decomposed into several 'stains' localized on a very short time-scale. In the same manner, the peaks, even of large amplitude, that do not correspond to stellar oscillations but are artefacts arising from the sampling or from the gaps of the observing window will appear in the time–frequency diagram localized on a very short time-scale, in a large frequency range (vertical atoms). A simple filtering keeping only the long-time-scale atoms then allows us to eliminate the information corresponding to noise (see Fig. 4 for the same diagram as in Fig. 3, but after filtering).

Since for a stable reconstruction we have to give up restoring the signal on to the whole frequency support, we are then led to select particular intervals over which the deconvolution makes sense. To address this aim, the most important groups of atoms are selected according to the criteria of high energy and sufficient lifetime. The oscillations with too short a lifetime correspond to a large band in the Fourier spectrum, and we cannot hope to restore the associated frequencies. If they are very energetic, they of course may produce perverse effects, but otherwise they only contribute to the background noise.

The peaks to be selected in the Fourier transform can now be precisely read on the y-axis of the matching pursuit diagram (see Fig. 4). Thus the role of the selection of the atoms is to detect the coherent structures lost in the noise, and then to choose on the

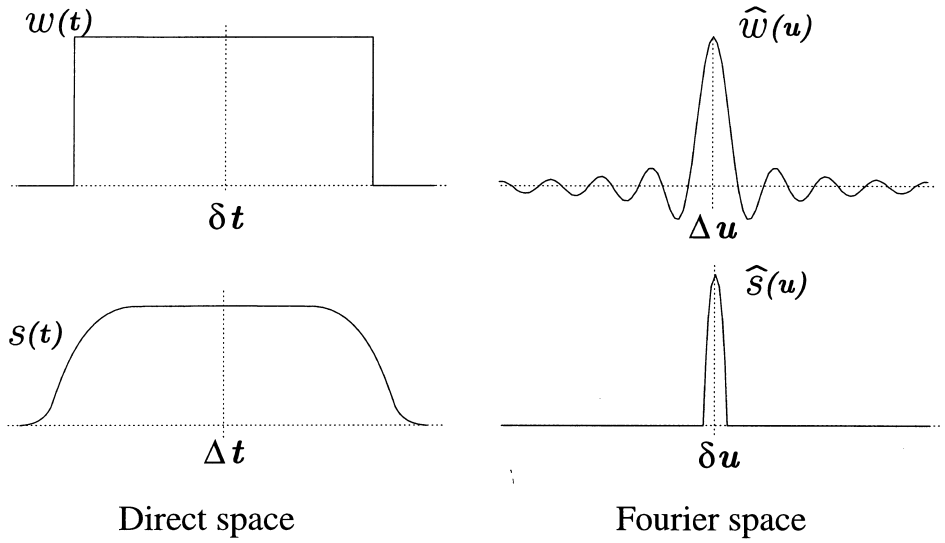

Figure 5. Regularization parameters in direct and Fourier space. The characteristic function of the observing window is $w(t)$ and the size of its support is δt points. Its Fourier transform is $\hat{w}(u)$. Its support is not bounded, but the term Δu represents the size of the central interval without sidelobes. The smoothing impulse function $\hat{s}(u)$ defines the target resolution on the restored spectrum and its support includes δu points. Its inverse Fourier transform $s(t)$ has 95 per cent of its energy in Δt (which can be identified with W_r of Fig. 2).

Fourier transform the frequency intervals over which the deconvolution has to be performed. Note that the successive selected peaks do not necessarily correspond to structures of strictly decreasing amplitude. Afterwards, the exact size of each interval is determined so that the border effects are minimized, for example by choosing the ends of the interval D_f on local minima of the frequency.

Once the deconvolution has been performed over one support, we can set in motion the deconvolution on to the following one, and so on. Under the condition that the spectrum of the star can be broken up, which is physically the case, we can reconstruct it step by step at a target gain in resolution, whereas a global reconstruction is impossible at the same resolution.

5 APPLICATION TO THE LIGHT CURVE OF GD 154

We now present the example of the DA variable white dwarf GD 154. In 1991 May, the star was the target of a WET campaign (Pfeiffer et al. 1996) (162.5 h of data collected). The $f_1 = 0.843$ mHz (1186 s) frequency of Robinson et al. (1978) was still present, but that at $1.52f_1$ had disappeared. Other frequencies ($f_2 = 0.918$ mHz and $f_3 = 2.484$ mHz) were also observed. Rotational splitting (with a period of 2.3 d) was detected in the f_1 peak. Pfeiffer et al. propose an identification of the modes (including $1.52f_1$) which yields a mass of the outer hydrogen layer of $2 \times 10^{-10} M_\star$ (where $M_\star$ is the mass of the star).

The time series that we have used in this work is extracted from this 1991 WET campaign; the corresponding light curve, acquired with the US Air Force 24-in telescope at Mauna Kea Observatory on 1991 May 17, and the corresponding power spectrum are presented in Fig. 1. The variations are multiply periodic, with periods in the range 100–3000 s. The general appearance of $\hat{\psi}$ shows that the high-frequency modes are more or less visible, and the corresponding high-resolution details may be due to noise. We purposely chose our light curve as a small part from multi-site observations, as it will allow us to present a comparison with the result of the whole WET campaign, as a test of the proposed method. We will also give the result of the same procedure applied to several other single-night runs extracted from the same 1991 observations.

5.1 Restoration and analysis from a one-night observation

The light curve consists of a series of evenly spaced data obtained by photon-counting integrated on 10-s intervals. The length of the observing window is approximately 7 h ($\delta t = 2620$ points – see Fig. 5). The data are extended so that the number of sampling points is a power of two, by adding zeros on each side of the record ($N = 16384$ points).

The signal-to-noise ratio (equation 4) is computed according to a soft-thresholding procedure (see Appendix A). The parameters α and β which occur in the definitions (6) and (8) of the functions $\tilde{\psi}(t)$ and $g(t)$ are taken to be equal to 1.1 and 3 respectively. This choice is not critical, but α must of course be of the order of unity to ensure that we have at least as much signal as noise.

If we first try to deconvolve the whole spectrum without frequency interval discrimination, we obtain the following results: for a gain in resolution of about 2, the interpolation parameter (equation 10) is of the order of 10. For such η, the minimum eigenvalue μ is of the order of 10^{-3}, leading to a relative reconstruction error of about 80 per cent. This case corresponds of course to an ill-conditioned situation, and either the size of W_r or that of D_f must be reduced as indicated for example in Section 4 (last paragraph).

Let us now apply to the same data the procedure with frequency interval discrimination. The effective size of the support of the impulse response $\hat{w}(t)$, Fourier transform of the observing window function is $\Delta u = 15$ points, and this function is normalized so that $\hat{w}(0) = 1$. The deconvolution is applied to the frequency support corresponding to the second time–frequency atom from the top of Fig. 4 (around 2.5 mHz). This support includes 31 points. The region W_r, regularizing the support of the observing window function W, is a time interval of width $\Delta t = 3819$ points. As for an

energy rate of $\chi = 0.95$ for $s(t)$ (see Fig. 2) one has $\delta u \Delta t \simeq 0.95$ (Lannes, Casanove & Roques 1987b), the corresponding frequency interval support of $\hat{s}(t)$ (equation 3) defining the resolution limit of the reconstruction process then includes seven points ($\delta u = 7$). The expected gain in resolution, given by the formula $\gamma_r = \Delta u / \delta u$ is then $\gamma_r = 2.1$. These aspects are outlined in Fig. 5.

For the gain chosen, the interpolation parameter η (equation 10) is of the order of 3.7. The corresponding approximation to the minimum eigenvalue of the imaging operator is $\mu = 0.083$, and the upper-bound Θ is of the order of 17 per cent. As the situation seems to be relatively stable, the reconstruction process can then be set in motion. In 10 iterations the method of conjugate gradients provides the least-squares solution presented in Fig. 6. With the gain in resolution equal to 2.1, we could hope to separate some peaks if necessary. Thus, once the result has been obtained, we observe additional information and in particular the splitting of the two main peaks.

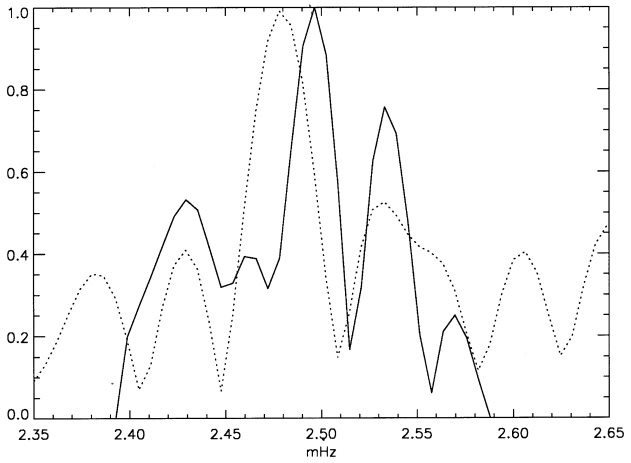

Figure 6. Solid line: normalized result of deconvolution on the support corresponding to a high-frequency structure given by the matching pursuit algorithm (precision: 0.02 mHz). Dotted line: raw Fourier transform.

For clarity, we now concentrate on the eigenvalue decomposition of the imaging operator. The error term of equation (12) corresponds to the case in which the overall error would be concentrated in the eigenspace associated with the minimum eigenvalue, which is not the case in practice. However, the minimum eigenvalue is the only one that we can reach easily before reconstruction. In fact, the error is distributed amongst all the eigenspaces. The conjugate gradient method allows us to compute the eigenvalues of A^*A and their associated weights κ_j. These terms represent the cosines of the 'separation angle' between the solution and the eigenvectors, and can give the critical eigenmodes: for example, for a small eigengenvalue, if κ_j is close to 1, the energy in the corresponding eigenmode is very large, so that the corresponding details may be artefacts. In such situations, the spectrum to be reconstructed may be redefined at a lower level of resolution, and the reconstruction process set in motion again. The first four eigenvalues μ_j and their corresponding weights κ_j are listed in Table 1.

The actual error corresponding to the example presented proves to be of the order of 13 per cent. As the results of this error analysis are satisfactory, the reconstruction process need not be be set in motion again.

5.2 Comparison with multi-site data

The choice of processing a one-night time series extracted from a

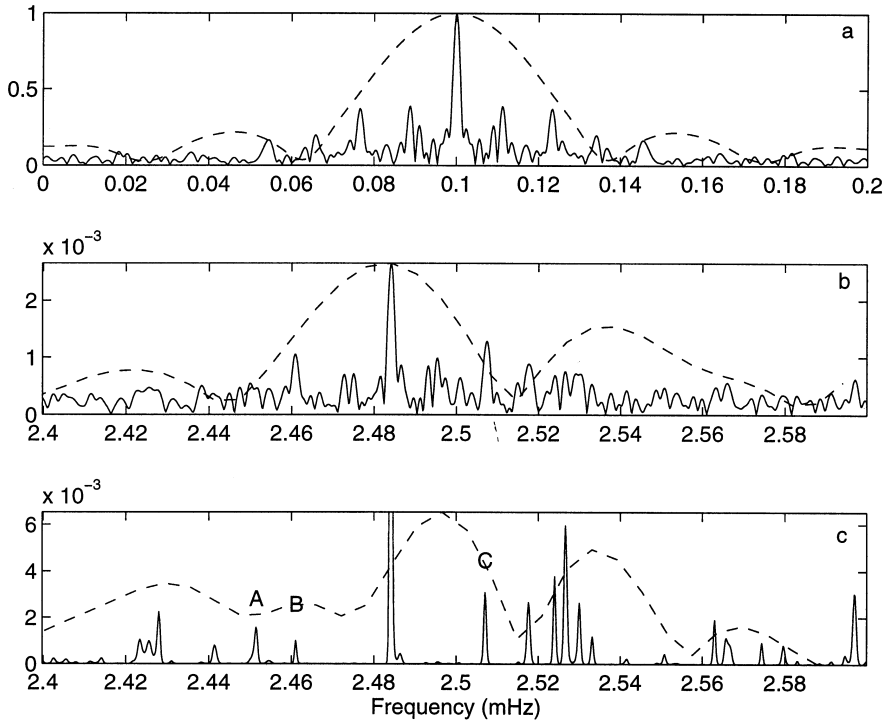

Figure 7. (a) WET (solid line) and one-night (dashed line) spectral windows. (b) WET (solid line) and one-night (dashed line) spectra in the frequency interval of interest. (c) WET (solid line) and one-night (dashed line) deconvolutions, using a Richardson–Lucy algorithm for the WET data, and our algorithm for the one-night data.

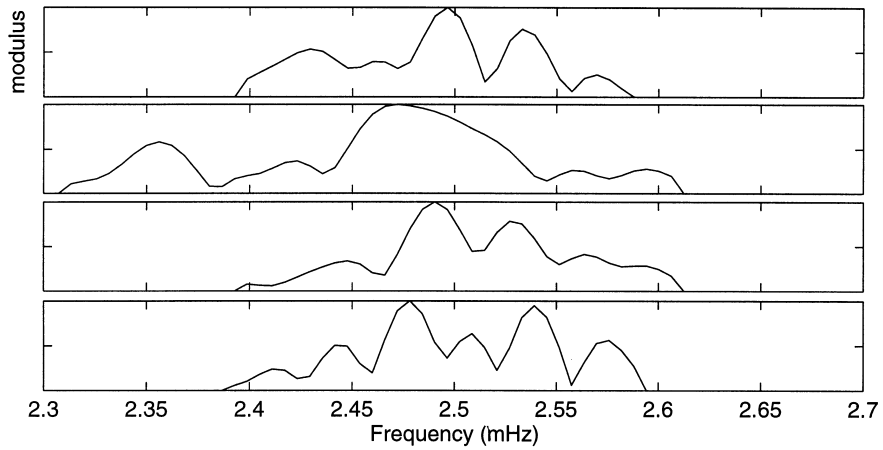

Figure 8. Four deconvolution results on different nights. First panel: deconvolution result previously discussed. Second panel: observation on 1991 May 16, on the same telescope. Third and fourth panels: runs obtained on 1991 May 22 and 24, at the CFH 3.6-m telescope.

Table 1. Eigenvalue decomposition. The first column represents the first four eigenvalues and the second column their respective weights.

μ_j	κ_j
0.083	0.0287
0.093	0.0144
0.410	0.0128
0.430	0.0060

WET campaign allows us now to analyse the results by comparison with those of complete multi-site observations. Of course, the frequency resolution achieved by the WET observations is much better than what we could expect from our reconstruction process. However, one can see (Fig. 7b) that in such multi-site data sets there is a strong noisy background, together with a number of spurious features of no physical signifiance, owing to the sidelobes of the impulse response (solid line in Fig 7a).

Fig. 7(b) presents the raw WET spectrum, together with the raw Fourier transform of the data that we have processed in this paper. Fig. 7(c) displays the result of a damped, low-pass-filtered Richardson-Lucy (RL) algorithm (Richardson 1972; Lucy 1974; Tittley 1998), superimposed on the result of our deconvolution. The RL algorithm, based on Bayesian inference, solves inverse problems by interpreting the kernel of the problem (here the spectral window) as a probability. The RL method finds much use in the field of deconvolution in astronomy. The RL algorithm does not give a perfect result here, since there is still some leakage from the sidelobes of the window function: peaks B and C and one of the peaks around 2.52 mHz in Fig. 7(c) are probably artefacts. This is due to the necessity of limiting the error propagation in the deconvolution (and hence the gain in resolution) because of the strong noisy background of the raw WET data in the regions of interest.

Nevertheless, Fig. 7(c) shows a good correlation between the peaks of our deconvolution and the RL results. The frequency difference between our main peak and the WET one at 2.484 mHz is well within the frequency resolution of the deconvolution (precision: 0.02 mHz). The other peaks near 2.42, 2.53 and 2.57 mHz correspond to regions where energy appears also in the RL spectrum. Note that the second of those peaks is also in good agreement with Pfeiffer et al. (1996), where two frequencies are listed at 2.5186 and 2.5262 mHz. The situation is not so clear for the peak at 2.46 mHz, which could correspond to peak A of the RL solution.

If we now compare with the raw Fourier transform of our data set (dashed line in Fig. 7b), we confirm that the gain in spectral resolution, allowing the separation of several peaks of the Fourier transform, brings really new valuable information. Moreover, we can see the elimination of the sidelobes of the one-night spectral window: its shape (dashed line in Fig. 7a) indicates clearly that the sidelobes are indistinguishable from, for example, a real peak near 2.54 mHz.

5.3 Comparison of several runs

To confirm the result of our single-night deconvolution, we applied the same procedure to three other nights of the same WET campaign. The resulting spectra are collected in Fig. 8 (the first panel recalls the result previously discussed).

The second panel (another night on the 24-in telescope at Mauna Kea) has a lower resolution, owing to the fact that the data set used was much shorter. The third and fourth panels display results from two nights at the Canada–France–Hawaii (CFH) 3.6-m telescope [see Pfeiffer et al. (1996) for a description of those three runs]. From these figures we can conclude the following.

(i) The main peak is not obtained at exactly the same frequency, but the differences are within the spectral resolution of the deconvolution.

(ii) The peak around 2.53–2.54 mHz is well-confirmed (even if not well-separated from the main peak in the second panel, owing to the lack of resolution). A peak near 2.57 mHz, with a smaller amplitude, seems to be present in the four runs, in good agreement with the result displayed in Fig. 7(c) for the whole campaingn. Around 2.42–2.46 mHz, peaks of low amplitude are quite variable

from one run to the other, and any interpretation would be hazardous.

6 CONCLUSION

During the last 10 years, the observational tools used for asteroseismology of rapidly oscillating stars have been greatly improved, mostly owing to the setting up of the Whole Earth Telescope (Nather et al. 1990), which allows a much better temporal coverage than one-site observations. On the other hand, important theoretical advances have been made at the same time, with the development of very accurate grids of white dwarf models and oscillation calculations (Brassard et al. 1992; Bradley 1996).

At this stage, it is also essential to improve the signal analysis, in order to get the maximum information from the observations: the standard Fourier analysis, along with traditional reconstruction methods used in asteroseismology, such as 'pre-whitening' (Ponman 1981) or CLEAN (Schwarz 1978), suffer from severe limitations. The main problems with these methods are: first, they may introduce some bias, especially when there are interactions between peaks of the Fourier transform; secondly, they are not regularized, and as a result the condition number (the ratio of the maximum and the minimum eigenvalues of the imaging operator) of the algorithms may be very large, producing an over-resolution of the solutions.

This paper has been an attempt to set out the main guidelines of a unified approach to the reconstruction problem encountered in asteroseismology. We are not attempting at this stage to present new astrophysical results, but only to demonstate the power of more advanced methods of signal analysis in this domain.

The deconvolution problem is stated in terms of weighted interpolation with partial extrapolation. In the case of time series with no gap, we demonstrate the possibility of obtaining better spectral resolution than the classical Fourier transform on a limited spectral interval.

Thanks to time–frequency analysis, we have proposed a process that allows more precise and selective time-series analysis. This process is rich enough to represent efficiently all coherent structures of the signal: the matching pursuit algorithm makes use of a dictionary from which the atoms are selected, and which is much larger than a basis. In the future of rapidly oscillating star asteroseismology, more and more attention will be paid to the physics of the non-linear behaviour of the stars, which can give rise to modulation of the frequencies and amplitudes of the oscillations. The introduction of time–frequency analysis, as in the matching pursuit algorithm, will prove to be a very useful tool in this domain, but it will be necessary to add in the dictionary atoms depending on the particular nature of this type of signal.

This way of presenting concrete quantitative results is the first step towards deconvolution of multi-site observations, and allows us to characterize the artefacts generated by other restoration techniques. In a forthcoming work, we will carry out more investigations in the problem of gaps and overlapping in the WET records by means of parametric spectral analysis.

REFERENCES

Anscombe F. J., 1949, Biometrics, p. 165
Bouyoucef K., Fraix-Burnet D., Roques S., 1997, A&AS, 121, 575
Bracewell R. N., 1955, J. Opt. Soc. Am., 45, 1756
Bradley P., 1996, ApJ, 468, 350
Brassard P., Fontaine G., Wesemael F., Hansen C. J., 1992, ApJS, 80, 369
Breger M., Handler G., Nather R. E., 1995, A&AS, 297, 473
Charpinet S., Fontaine G., Brassard P., Dorman B., 1996, ApJ, 471, 103
Clemens J. C., 1993, in Weigend A. S., Gershenfeld N. A., eds, SFI Studies in Science Complexity, Proc. Vol. XV, Time Series Predictions. Addison-Wesley, Reading, MA, p. 139
Cohen A., Daubechies I., Jawerth B., Vial P., 1993, C. R. Acad. Sci. Paris A, 316, 417
Daubechies I., 1992, Ten lectures on wavelets. Society for Industrial and Applied Mathematics, Philadelphia
Donoho D., 1994, in Meyer Y., Roques S., eds, Progress in Wavelet Analysis and Applications. Editions Frontières, Gif-sur-Yvette, p. 109
Donoho D., Johnstone I. M., 1994, Biometrika, 81, 425
Hestenes M., Stiefel E., 1952, J. Res. Natl. Bur. Stand., 49, 409
Kilkenny D., Koen C., O'Donogue D., Stobie R. S., 1997, MNRAS, 196, 583
Koester D., Chanmugan G., 1990, Rep. Prog. Phys., 53, 837
Lannes A., Roques S., Casanove M. J., 1987a, J. Mod. Opt., 2, 34, 161
Lannes A., Casanove M. J., Roques S., 1987b, J. Mod. Opt., 3, 34, 321
Lannes A., Roques S., Casanove M. J., 1987c, J. Opt. Soc. Am. A, 1, 4, 189
Lucy L. B., 1974, AJ, 79, 745
Mallat S., 1989, IEEE Trans. PAMI, 7, 11, 674
Mallat S., Zhang Z., 1993, IEEE Trans. Signal Process., 12, 41, 3397
Michel E. et al., 1992, A&A, 255, 139
Nather R., Winget D., Clemens J., Hansen C., Hine B., 1990, ApJ, 361, 309
Pfeiffer B. et al., 1996, A&A, 314, 182
Ponman T., 1981, MNRAS, 196, 583
Richardson W. H., 1972, J. Opt. Soc. Am., 62, 55
Robinson E. L., Stover R. J., Nather R. E., McGraw J. T., 1978, ApJ, 220, 614
Roques S., Bourzeix F., Bouyoucef K., 1996, Ap&SS, 239, 297
Sanz J. L. C., Huang T. S., 1983, J. Opt. Soc. Am., 11, 73, 1455
Schwarz U. J., 1978, A&A, 65, 345
Slepian D., 1964, Bell Syst. Tech. J., 43, 3009
Tikhonov A. N., Arsenin V. Y., 1977, Solutions of ill-posed problems. Winston, Washington DC
Tittley E. R., 1998, Richarson–Lucy software, URL: http://phobos.astro.uwo.ca/~etittley/projects
Ville J., 1948, Cables Transm., 2A, 1, 61
Wickerhauser M. V., Coifman R. R., 1992, IEEE Trans. Inf. Theory, 32, 712
Youla D. C., 1978, IEEE Trans. Circuits and Systems, CAS-25, 694

APPENDIX A: DENOISING BY SOFT THRESHOLDING

A natural way of separating the noise from a signal is to use a 'multiresolution analysis', which provides a representation intermediate between the spatial and the Fourier ones: the data are represented as a superposition of wavelets at some scale level, each level being further decomposed into a lower scale level. The reader unfamiliar with this analysis is referred to Mallat (1989) or Daubechies (1992).

Using orthogonal bases of compactly supported wavelets, a soft-thresholding noise removal method (Donoho & Johnstone 1994; Donoho 1994) can be applied in the wavelet domain. It consists of non-linearly shrinking the empirical wavelet coefficients. This allows us to recover the noise-free signal, and then the noise. This noise, compared with raw data, yields the pointwise signal-to-noise ratio (equation 4).

More precisely, if ψ is the noisy signal, the algorithm proposes a three-step method for recovery of the noise-free signal.

(i) Apply the pyramidal wavelet filtering of Cohen et al. (1993) to the new data $\psi/\sqrt{N}$ (N is the number of sampling points), yielding noisy wavelet coefficients z.

(ii) Apply the soft-thresholding $\zeta_\tau(z) = \mathrm{sgn}(z)(\sup(0, |z| - \tau)$ to the noisy wavelet coefficients z, with threshold $\tau = \sqrt{2 \log N}\sigma/\sqrt{N}$, where σ^2 is the variance of the noise estimated in the data. This operation shrinks by the amount τ the wavelet coefficients towards zero, and keeps only the non-negative elements. With a photon-counting model, it is possible first to apply the Anscombe variance-stabilizing transform (1949) point by point to the noisy data ψ, and then to act as if the data arose from a Gaussian white noise model of variance unity.

(iii) Invert the wavelet transform, producing the noise-free signal.

One can then recover the noise by subtraction of the noise-free signal from the noisy data. The advantages of this operation are presented in an astrophysical framework by Roques, Bourzeix & Bouyoucef (1996).

This paper has been typeset from a TEX/LATEX file prepared by the author.