Ultrasonic non destructive testing based on sparse deconvolution

To cite this article: C Soussen *et al* 2012 *J. Phys.: Conf. Ser.* **353** 012018

View the article online for updates and enhancements.

Ultrasonic non destructive testing based on sparse deconvolution

C Soussen[1], J Idier[2], E Carcreff[2,3], L Simon[3] and C Potel[3,4]

[1] Centre de Recherche en Automatique de Nancy (CRAN, UMR 7039, Nancy-University, CNRS). Campus Sciences, B.P. 70239, F-54506 Vandœuvre-lès-Nancy, France. This work has been performed while the first author was a visiting research fellow at IRCCyN, appointed by the CNRS

[2] Institut de Recherche en Communications et Cybernétique de Nantes (IRCCyN, UMR CNRS 6597), BP 92101, 1 rue de la Noë, F-44321 Nantes Cedex 3, France

[3] Laboratoire d'Acoustique de l'Université du Maine (LAUM, UMR CNRS 6613). Avenue O. Messiaen, F-72085 Le Mans Cedex 9, France

[4] Fédération Acoustique du Nord-Ouest (FANO, FR CNRS 3110), France

E-mail: `Charles.Soussen@cran.uhp-nancy.fr`

Abstract. The acoustic modality yields non destructive testing techniques of choice for in-depth investigation. Given a precise model of acoustic wave propagation in materials of possibly complex structures, acoustical imaging amounts to the so-called acoustic wave inversion. A less ambitious approach consists in processing pulse-echo data (typically, A- or B-scans) to detect localised echoes with the maximum temporal (and lateral) precision. This is a resolution enhancement problem, and more precisely a sparse deconvolution problem which is naturally addressed in the inversion framework. The paper focuses on the main sparse deconvolution methods and algorithms, with a view to apply them to ultrasonic non-destructive testing.

1. Introduction

Non-destructive testing by pulse-echo ultrasound allows the inspection of a complex object and in particular, the detection and localization of inhomogeneities and faults. The acoustic signal measured by the ultrasonic sensor results from the propagation inside the inspected object, of the waves generated by the interaction of the incident wave with the object. Reflections and/or transmissions of the wave induce the presence of echoes of the incident wave in the measured signal. Echoes are closely related to the change of acoustic impedance inside the object. Their detection and localization in terms of time of arrival and relative amplitude can be formulated as a spike deconvolution problem where the impulse response corresponds to the acoustic wave emitted from the input transducer and the spike locations and amplitudes are related to the acoustic reflectivity of the medium. Therefore, their accurate estimation is needed to obtain quantitative information about the inner structure of the object, both in terms of geometry and of structure.

From the information processing viewpoint, the spike deconvolution problem is especially challenging when the material dimension is small in comparison with the wavelength of the incident signal because the echoes overlap. In such a case, the discrimination and precise localisation of the echoes is far from obvious. Another issue occurs when the noise level is

high since echoes may be partly or completely drowned in noise due to the signal attenuation while the wave is propagating.

The simplest method to detect and localize echoes is the matched filter [1]. It is based on the maximization of the inter-correlation between the measured signal y and the incident wave h. Let us remark that matched filtering lends itself to a simple interpretation in terms of least-squares principle. It simply consists in considering an isolated echo of unknown delay τ and relative scalar amplitude x, and of minimising a least-squares fidelity-to-data criterion with respect to both x and τ. Since the criterion is a quadratic function of the amplitude x, the minimisation can be carried out explicitly with respect to x, and the remaining minimisation step with respect to τ amounts to maximising the inter-correlation between y and h.

Unfortunately, this technique does not yield satisfying results when the echoes overlap. The approach addressing the deconvolution problem using the assumption that the deconvolved signal is *sparse* (there are a very limited number of spikes) is known to yield better results in that the time resolution of the above mentioned approach can be significantly enhanced [2]. In particular, the approach which models the reflectivity sequence using the Bernoulli-Gaussian model [3–5] is acknowledged to be very efficient. It was successfully applied to the deconvolution of ultrasonic echoes in composite materials [6].

In this paper, we address the spike deconvolution as an inverse problem where the detection and estimation of the spikes are performed jointly. The inversion method explicitly takes into account that the reflectivity signal is sparse. This can be done using either ℓ_1 [7] or ℓ_0 regularisation [8]. In particular, the latter approach has close connection with the Bernoulli-Gaussian deconvolution approach [9]. In both cases, the inversion problem results in a numerical optimisation problem. In the ℓ_1 case, the optimisation problem is continuous and recent efficient algorithms are available [10]. For ℓ_0 regularisation, the optimisation problem turns into a discrete search problem whose goal is to find the spiky configuration (with a limited number of spikes) yielding the best fit to the data [11].

The paper is organised as follows. In Section 2, we formulate the sparse deconvolution problem. In Section 3, we introduce the ℓ_0 and ℓ_1 deconvolution algorithms and we address the estimation of the number of spikes from the algorithm outputs using classical order model selection rules [12]. Section 4 illustrates our approach on two simulated examples: the first one is a classical spike train deconvolution example introduced in the seismic reflection context in [3]. It will enable us to compare ℓ_0 and ℓ_1 deconvolution methods. The second example simulates a more realistic situation of the propagation of an acoustic wave inside a homogeneous material of known dimension and structure. This provides an accurate analysis of the benefit of recent sparse deconvolution for ultrasonic data analysis.

2. Sparse deconvolution as an inverse problem
2.1. Definition of the convolutive model
We use the convolutive observation model

$$y(t) = \sum_{k=1}^{K} x^\star(t_k)\, h(t-t_k) + \varepsilon(t)$$

to describe the propagation of the ultrasonic wave inside a material. h stands for the incident acoustic wave delivered by the input transducer, and t_k and $x^\star(t_k)$ model the time of arrival and the amplitude attenuation corresponding to the k-th echo. The noise $\varepsilon(t)$ includes the measurement noise due to the ultrasonic sensor and the modeling errors. Because the

measurements are discrete, we resort to the more realistic formulation

$$y(mT) = \sum_{k=1}^{K} x^\star(t_k)\, h(mT - t_k) + \varepsilon(mT),\ m = 0, \ldots, M-1 \tag{1}$$

where T is the sampling period and M is the number of measurements.

At this point, it is important to notice that t_k are continuous valued. For simplicity reasons, we choose to sample t_k over a grid with sampling period T. This leads to the discrete convolutive model

$$y(mT) = \sum_{n=0}^{N-1} x^\star(nT)\, h((m-n)T) + \varepsilon(mT),\ m = 0, \ldots, M-1 \tag{2}$$

where $x^\star(nT)$ is zero valued for most n, except when nT coincides with the time of arrival of an echo (*i.e.*, a t_k value in (1)). In the following, we use vector notations $\boldsymbol{y} \in \mathbb{R}^M$, $\boldsymbol{h} \in \mathbb{R}^P$ and $\boldsymbol{x}^\star \in \mathbb{R}^N$ to refer to the discrete signals resulting from the sampling of $y(t)$, $h(t)$ and $x(t)$. Because the convolutive model (1) is causal, we set $M = N + P - 1$ so that any echo resulting from $\boldsymbol{x}^\star$ is fully observable in signal $\boldsymbol{y}$. In matrix form, model (2) reads

$$\boldsymbol{y} = \boldsymbol{h} * \boldsymbol{x}^\star + \boldsymbol{\varepsilon} \tag{3}$$

where $*$ stands for the discrete convolution operator.

2.2. Inversion of the convolutive model

The sparse deconvolution problem consists in estimating the signal $\boldsymbol{x}^\star$ knowing $\boldsymbol{h}$ and the observation $\boldsymbol{y}$. In practice, $\boldsymbol{h}$ is the emitted ultrasonic wave and $\boldsymbol{y}$ is the recorded signal. For the sake of clarity, $\boldsymbol{x}^\star$ denotes the exact (unknown) sparse signal and we will denote by $\boldsymbol{x}$ an estimated version of $\boldsymbol{x}^\star$ deduced from $\boldsymbol{y}$ and $\boldsymbol{h}$.

As many other inverse problems, our deconvolution problem is *ill-posed* since $\boldsymbol{h}$ plays the role of a filter of limited bandpass. This means that there exist many signals $\boldsymbol{x}$ such that $\boldsymbol{h} * \boldsymbol{x}$ is close to the measured data $\boldsymbol{y}$, some of them being very different from the true solution $\boldsymbol{x}^\star$. The least-squares estimation of $\boldsymbol{x}^\star$ consists in minimising the criterion $\|\boldsymbol{y} - \boldsymbol{h} * \boldsymbol{x}\|^2$, where $\|\,.\,\|^2$ stands for the usual ℓ_2 norm defined as the sum of squares. However, this basic approach ignores the ill-posed character of the deconvolution problems. It provides solutions that are highly sensitive to the noise embedded in the data [2]. To obtain more stable solutions, it is necessary to introduce appropriate *regularisation* by imposing some constraints on $\boldsymbol{x}$ [13].

In the present context, it is very natural to impose a sparsity constraint on $\boldsymbol{x}$, *i.e.*, that $\boldsymbol{x}$ has only a few non-zero values. This can be done by solving the constrained least-squares problem

$$\min_{\boldsymbol{x}} \|\boldsymbol{y} - \boldsymbol{h} * \boldsymbol{x}\|^2 \quad \text{subject to} \quad \sum_n \varphi(x_n) \leqslant \lambda, \tag{4}$$

where the constraint is designed to favor sparse solutions and the parameter λ controls the compromise between fidelity to data and the level of sparsity. Several choices are available for φ. In the following, we focus on ℓ_0 and ℓ_1 regularisations.

2.3. ℓ_0 deconvolution

The ℓ_0 cost function is defined in one dimension as

$$|t|_0 = \begin{cases} 0 \text{ if } t = 0, \\ 1 \text{ otherwise.} \end{cases} \tag{5}$$

By extension, the ℓ_0 pseudo-norm of a vector $\boldsymbol{x}$ is defined as $\|\boldsymbol{x}\|_0 = \sum_m |x_m|_0$. It counts the number of non-zero elements in $\boldsymbol{x}$.

When setting $\varphi(t) = |t|_0$, problem (4) equivalently reads

$$\min_{\boldsymbol{x}} \|\boldsymbol{y} - \boldsymbol{h} * \boldsymbol{x}\|^2 \quad \text{subject to} \quad \|\boldsymbol{x}\|_0 \leqslant K, \tag{6}$$

where $\lambda = K$ is the maximum number of spikes allowed in the signal $\boldsymbol{x}$. (6) is actually a combinatorial optimisation problem over $\{0,1\}^M$ because it consists in finding the support of $\boldsymbol{x}$ of dimension K yielding the least residual error, and there are a finite but large number of configurations with K spikes. Unfortunately, such a problem cannot be solved exactly unless (nearly) all solutions are tested. In practice, ℓ_0 deconvolution methods only explore a subset of all solutions, on the basis of heuristic rules. They are thus *suboptimal*, in the sense that they are not guaranteed to find the true solution of (6).

2.4. ℓ_1 deconvolution

In the sparse signal approximation literature, regularisation using the ℓ_1 norm has become very popular because it leads to convex optimisation problems that can be exactly solved in contrast to their ℓ_0 counterparts, while also yielding sparse solutions.

The ℓ_1 norm is defined as $\|\boldsymbol{x}\|_1 = \sum_m |x_m|$, so the ℓ_1 deconvolution problem can be formulated as

$$\min_{\boldsymbol{x}} \|\boldsymbol{y} - \boldsymbol{h} * \boldsymbol{x}\|^2 \quad \text{subject to} \quad \sum_m |x_m| \leqslant \lambda. \tag{7}$$

The regularisation parameter λ does not stand for the number of spikes anymore. It represents the upper bound on the sum of amplitudes of the signal $\boldsymbol{x}$. However, the lower λ, the larger the number K of spikes in $\boldsymbol{x}$, so that there is an implicit relation between λ and K.

3. Algorithms

This section firstly introduces the algorithms we use to address the ℓ_0 and ℓ_1 deconvolution problems. Their outputs are denoted by $\boldsymbol{x}^0(K)$ and $\boldsymbol{x}^1(\lambda)$, respectively. In a second part, the issue of automatic selection of the regularisation parameter is dealt with, which amounts to determining the model order $K = \|\boldsymbol{x}\|_0$. Our approach can address both ℓ_0 and ℓ_1 cases.

3.1. ℓ_0 deconvolution algorithm

For reasons that will become clear in Subsection 3.3, we preferentially consider *greedy* iterative algorithms to address (6). Such numerical schemes provide a set of solutions $\boldsymbol{x}^0(K)$ for all cardinalities $K = 0, 1, \ldots, K_{\max}$ where $K_{\max}$ is a predefined maximal cardinality. *Forward* algorithms start from an empty support and then incrementally introduce new elements (*i.e.,* new spikes) in the support and re-estimate the amplitudes $\boldsymbol{x}$ whenever the support is modified [14, 15].

In deconvolution problems where many echoes frequently overlap, the forward approach is not fully satisfactory: typically, in the presence of two overlapping echoes of similar amplitudes, it is likely to firstly position a false spike in between the two true positions, and then to compensate this wrong initial guess by positioning additional false spikes.

Therefore, we rather favor *forward-backward* algorithms, *i.e.,* strategies that enable both the insertion and the removal of a spike at each iteration. A spike insertion or removal will be termed as a *single replacement*. In the following, we use the Single Best Replacement algorithm (SBR) algorithm to perform ℓ_0 deconvolution. The SBR algorithm is of the *forward-backward* type [9]. It is an extension of the Single Most Likely Replacement (SMLR) algorithm proposed in the

context of Bernoulli-Gaussian deconvolution [3–6, 16]. The SMLR algorithm is acknowledged to be very efficient for sparse deconvolution. The difference between SBR and SMLR is that SBR addresses the problem (6) while SMLR is dedicated to a slightly different problem where the regularisation also includes a penalty on the range of non-zero amplitudes x_m.

3.2. ℓ_1 deconvolution algorithm

Although the regularisation parameter λ is continuous valued, ℓ_1 regularisation is very attractive because efficient algorithms are available to compute the solution $\boldsymbol{x}^1(\lambda)$ of (7) for a continuous range of λ values ($\lambda \in [\lambda_{\min}, +\infty)$). The principle of the *homotopy* algorithm is to gradually decrease λ from $+\infty$ to $\lambda_{\min}$ by recursively computing the critical λ values for which the support of $\boldsymbol{x}^1(\lambda)$ is changing. For any critical λ-value, a series of single spike replacements are performed similar to the single replacements occurring in the forward-backward algorithms introduced above [17, 18]. Up to the storage of the intermediate spike replacements occurring during the homotopy iterations, homotopy not only yields a sequence of solutions $\boldsymbol{x}^1(\lambda)$ over the continuous range $\lambda \in [\lambda_{\min}, +\infty)$, but also a series of sparse approximations $\boldsymbol{x}^1(K)$ for all cardinalities $K = 0, 1, \ldots, K_{\max}$ where $K_{\max}$ is a predefined maximum cardinality.

3.3. Model order selection

Model order selection is a well known problem for which many statistical and heuristic rules have been proposed. We refer the reader to [19] for a review of the main information criterion rules and to [12] for their application and adaptation to sparse approximation problems.

In the sparse deconvolution context, model order selection amounts to choose the cardinality K (*i.e.,* the ℓ_0 norm of the solution) which corresponds to selecting a presumed "most acceptable" trade-off between a small least-squares criterion and a high degree of sparsity. Here, we propose to take advantage from the fact that both SBR and homotopy algorithms naturally yield a series of sparse solutions, respectively $\boldsymbol{x}^0(K)$ and $\boldsymbol{x}^1(K)$, for all cardinalities between 0 and $K_{\max}$. Therefore, the best trade-off can be found by direct inspection of $J(\boldsymbol{x}^0(K))$ or $J(\boldsymbol{x}^1(K))$, $K = 0, 1, \ldots, K_{\max}$, as soon as an order selection criterion J has been selected.

Two main categories can be distinguished among the most popular rules for order selection.

The first category proceeds by minimising a compound criterion that directly balances fidelity to data and cardinality. This criterion takes the general form

$$\arg\min_K \left\{ J(K) = M \log(\|\boldsymbol{y} - \boldsymbol{h} * \boldsymbol{x}(K)\|^2) + \alpha K \right\} \tag{8}$$

where M stands for the number of measurements (the size of $\boldsymbol{y}$), $\|\boldsymbol{y} - \boldsymbol{h} * \boldsymbol{x}(K)\|^2$ is the error residual obtained with the solution of cardinality K and α is a constant. α is set to 2, $\log M$ and $2 \log \log M$ for the Akaike Information Criterion (AIC) [20], the Minimum Description Length (MDL) [21] and the Hannan and Quinn Criterion [22], respectively.

The second category is the family of cross validation methods (cross validation is also called PRESS statistic for Predicted REsidual Sums of Squares). Their principle is to partition the data $\boldsymbol{y}$ into two complementary subsets, *e.g.,* by leaving out a single observation y_m leading to a reduced observation signal $\boldsymbol{y}^{[m]}$ of size $M - 1$. Cross validation consists in testing whether the estimation $\boldsymbol{x}(\boldsymbol{y}^{[m]}, K)$ of $\boldsymbol{x}$ knowing the reduced data signal using the sparse approximation algorithm leads to an accurate prediction of observation y_m [23]. Formally, cross validation aims at minimising a criterion of the form

$$J(K) = \frac{1}{M} \sum_m \left(y_m - \hat{y}_m^{[m]}(K) \right)^2 \tag{9}$$

where $\hat{y}_m^{[m]}(K)$ stands for the m-th entry of $\boldsymbol{h} * \boldsymbol{x}(\boldsymbol{y}^{[m]}, K)$. The fact that cross validation effectively penalises high-order models can be explained by the phenomenon of *overfitting*: a

solution of excessive order will be more sensitive to noise, so it will not produce a good cross validation score. Generalised cross validation is an improved version where the prediction errors $(y_m - \hat{y}_m^{[m]}(K))^2$ are weighted by a factor depending on K [24]. The special case of sparse approximation algorithms leads to simplified expressions and computation of the cross validation criteria. The reader interested in further details is referred to the book chapters [25, Chapter 5] and [12].

In the simulations below, we compare several automatic order selection rules for the sparse deconvolution problem. We obtain that most of them lead to an over-estimation of the number of spikes and that the MDL criterion yields the best results. For further information regarding MDL, the reader is referred to the comprehensive references [26, 27].

4. Numerical simulations
The goal of the presented simulations is to illustrate that ℓ_0 and ℓ_1 deconvolution algorithms are well suited to acoustic deconvolution problems and to compare their performance. Comparisons are made in terms of ability to detect the presence of spikes and to perform their localisation in a precise manner. The practical efficiency of the order selection criteria introduced above are also compared in order to automatically evaluate the number of spikes. We first consider a classical spike deconvolution problem and then another problem which is typical of acoustic deconvolution.

4.1. Classical spike deconvolution problem
This simulation is close to that introduced in [3]. It is designed in order to analyse in what respect a given algorithm is able to discriminate close echoes and to detect echoes of small amplitudes which are partly drown either in noise of in larger echoes. Figures 1(a,b) display the wavelet $\boldsymbol{h}$ of size $P = 21$ samples, and the simulated data $\boldsymbol{y} = \boldsymbol{h} * \boldsymbol{x}^\star + \boldsymbol{\varepsilon}$ corresponding to $\|\boldsymbol{x}^\star\|_0 = 19$ spikes. The signal-to-noise ratio (SNR) is set to 20 dB. Obviously, there are very close spikes which result in overlapping echoes.

The outputs of the ℓ_0 and ℓ_1 deconvolution algorithms are displayed on figures 1(c,d), respectively, where the number of spikes has been estimated according to the MDL selection criterion. Both algorithms perform accurate detection and localisation of the main spikes. They clearly outperform the algorithms based on the computation of the inter-correlation between $\boldsymbol{y}$ and $\boldsymbol{h}$ such as matched filtering [2]. The zoom in of figures 1(e,f) shows that ℓ_0 deconvolution yields the most accurate results since the spike localisation is obviously better and the signal approximation $\boldsymbol{y} \approx \boldsymbol{h} * \boldsymbol{x}$ is more precise than the approximation obtained with ℓ_1 deconvolution. Figure 2 illustrates the performance of both deconvolution algorithms for all sparsity levels K. Both curves represent the approximation error $\|\boldsymbol{y} - \boldsymbol{h} * \boldsymbol{x}(K)\|^2$ versus the cardinality K. The ℓ_0 curve always lays below the ℓ_1 curve, therefore indicating that ℓ_0 deconvolution leads to better data approximations for any cardinality. Moreover, it is interesting to notice that the ℓ_0 curve lays below the star corresponding to the true solution $\boldsymbol{x}^\star$. This indicates that the ℓ_0 solution of cardinality 19 leads to a residual error which is lower than the noise variance.

In table 1, we test classical order selection criteria. The general trend is that AIC and cross validation tend to over-estimate the number of spikes. For the simulation of figure 1, MDL provides the most realistic order selection since it correctly estimates the true number of spikes (19) for ℓ_0 deconvolution and it yields the number which is the closest to the true number for ℓ_1 deconvolution. For higher noise levels (SNR = 0 dB), MDL under-estimates the number of spikes. This makes sense because the spikes of small amplitudes are totally drowned in noise, justifying that sparse deconvolution algorithms cannot detect them.

Because ℓ_0 deconvolution yields the best results, we set apart ℓ_1 deconvolution and we only consider the ℓ_0 algorithm in the second simulation below.
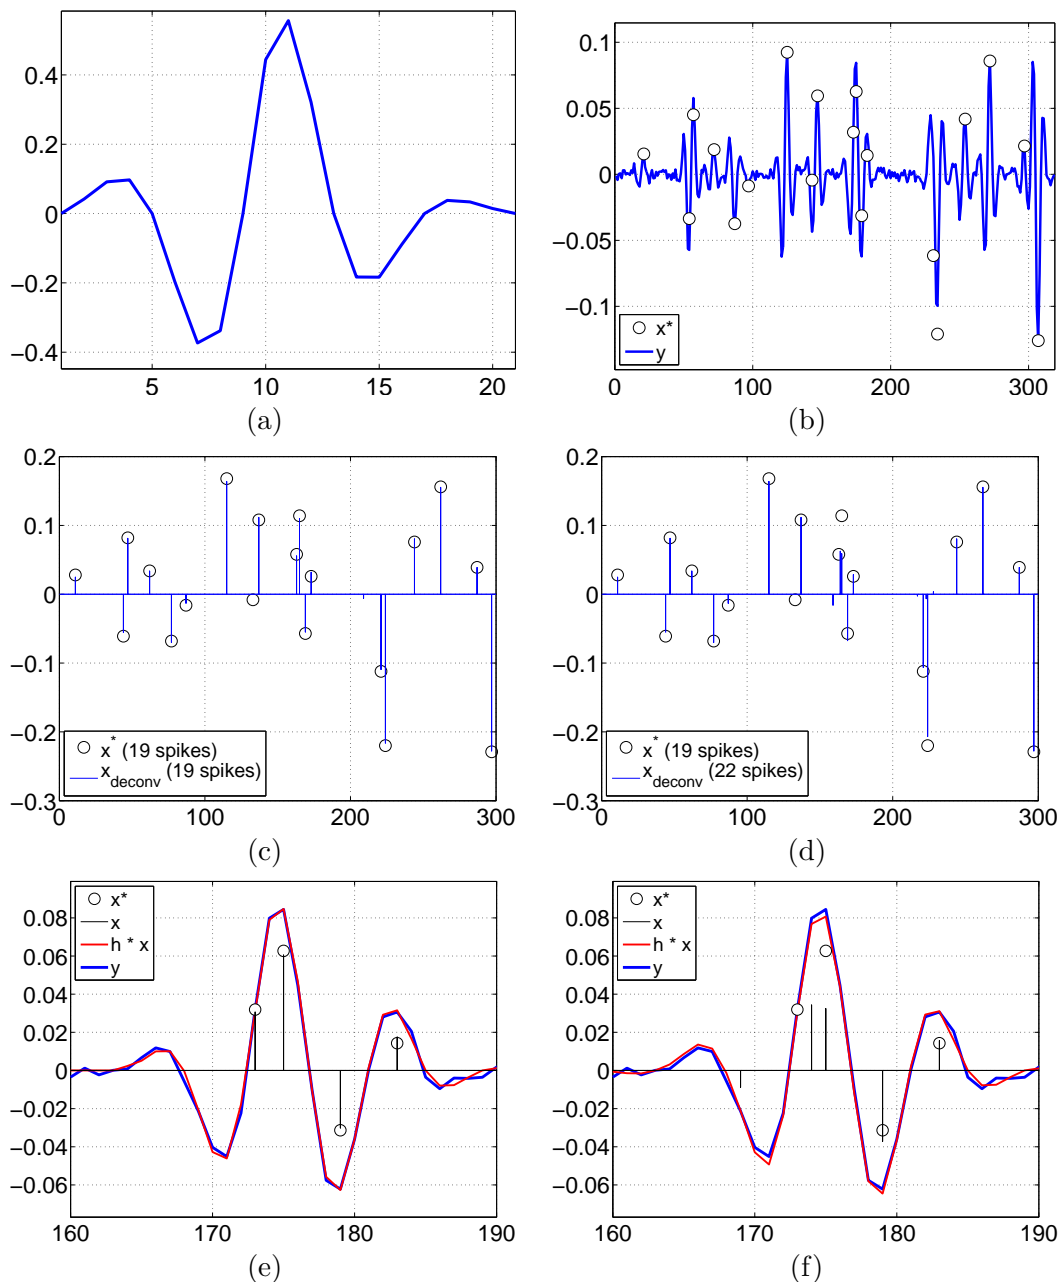
Figure 1. First simulation. (a) Wavelet signal $\boldsymbol{h}$, of size 21. (b) Display of the simulated data $\boldsymbol{y} = \boldsymbol{h} * \boldsymbol{x}^{\star} + \boldsymbol{\varepsilon}$ together with the true sparse signal $\boldsymbol{x}^{\star}$ ($\circ$). $\boldsymbol{y}$ and $\boldsymbol{x}^{\star}$ are of size 320 and 300, respectively. $\boldsymbol{x}^{\star}$ is composed of 19 spikes. The SNR is set to 20 dB. (c) Recovery of $\boldsymbol{x}$ using the ℓ_0 deconvolution method. (d) Recovery of $\boldsymbol{x}$ using the ℓ_1 deconvolution method. (e,f) Data approximations $\boldsymbol{y} \approx \boldsymbol{h} * \boldsymbol{x}$ (zoom in) using the outputs of the ℓ_0 and ℓ_1 deconvolution algorithms, respectively. The number of spikes is automatically estimated using the MDL criterion.

4.2. Processing of simulated acoustic data

This simulation is more specifically dedicated to the acoustic wave propagation problem. We simulate the propagation of the wavelet signal shown on figure 3 through an homogeneous isotropic aluminium material of known geometry according to the propagation model described

Figure 2. Complement to the simulation of figure 1: curves showing the outputs of the ℓ_0 and ℓ_1 deconvolution algorithms in terms of compromise between approximation error and cardinality. The ℓ_0 and ℓ_1 solutions (labeled ■ and ●) are obtained according to the MDL order selection criterion. The star label represents the cardinality of the true (unknown) solution $\boldsymbol{x}^\star$ and the related approximation error $\|\boldsymbol{h} * \boldsymbol{x}^\star\|^2$ corresponding to the noise power.

Table 1. Comparison of classical order selection criteria: Cross Validation (CV), Generalised Cross Validation (GCV), Akaike Information Criterion (AIC, AICs), Minimum Description Length (MDL, MDLs), Hannan-Quinn information Criterion (HQC). AICs and MDLs are modified versions of AIC and MDL criteria. They were derived to improve AIC and MDL when the number of observations is relatively small [28]. Tests are done on the simulation of figure 1 with moderate noise level (SNR = 20) and on a similar simulation with higher noise level (0 dB).

	CV	GCV	AIC	AICs	**MDL**	MDLs	HQC
SNR = 20 dB							
ℓ_0 deconvolution	41	39	39	28	**19**	18	26
ℓ_1 deconvolution	31	31	31	23	**22**	22	22
SNR = 0 dB							
ℓ_0 deconvolution	45	42	42	21	**11**	9	17
ℓ_1 deconvolution	29	30	30	15	**14**	14	15

in [29]. Obviously, the output signal resulting from various reflections of the wave at the edges of the material, has a content which highly depends on the incident angle and the dimension of the block of aluminium. We consider the case where the incident angle is normal to the material surface and the material thickness (2 mm) is smaller than the longitudinal wavelength (2.92 mm), so that the echoes strongly overlap in the data signal. Figures 4(a,c) display the data simulated for two different noise levels (SNR = 20 and 10 dB). It is noticeable that the echo amplitudes are decreasing with time due to the signal attenuation while the wave is propagating through the material. The ℓ_0 deconvolution method is able to perfectly detect and localise most echoes for SNR = 20 dB without any false detection. For a higher level of noise (SNR = 10 dB), the first main echoes are correctly localised but possible false detections are likely to occur for the smallest echoes which are drown in noise. The MDL selection criterion under-estimates the number of spikes (17 and 19 spikes are found instead of 22) since the smallest spikes cannot be detected. We believe that it is always desirable to under-estimate the number of spikes rather than detecting more spikes together with false alarms.

Finally, let us stress that the presented deconvolution methods are not computationally demanding. For instance, the results of figures 1 and 4 have been obtained in about 0.25

Figure 3. Realistic acoustic wavelet of size 100. It is a typical ultrasound signal propagating through an aluminium object.
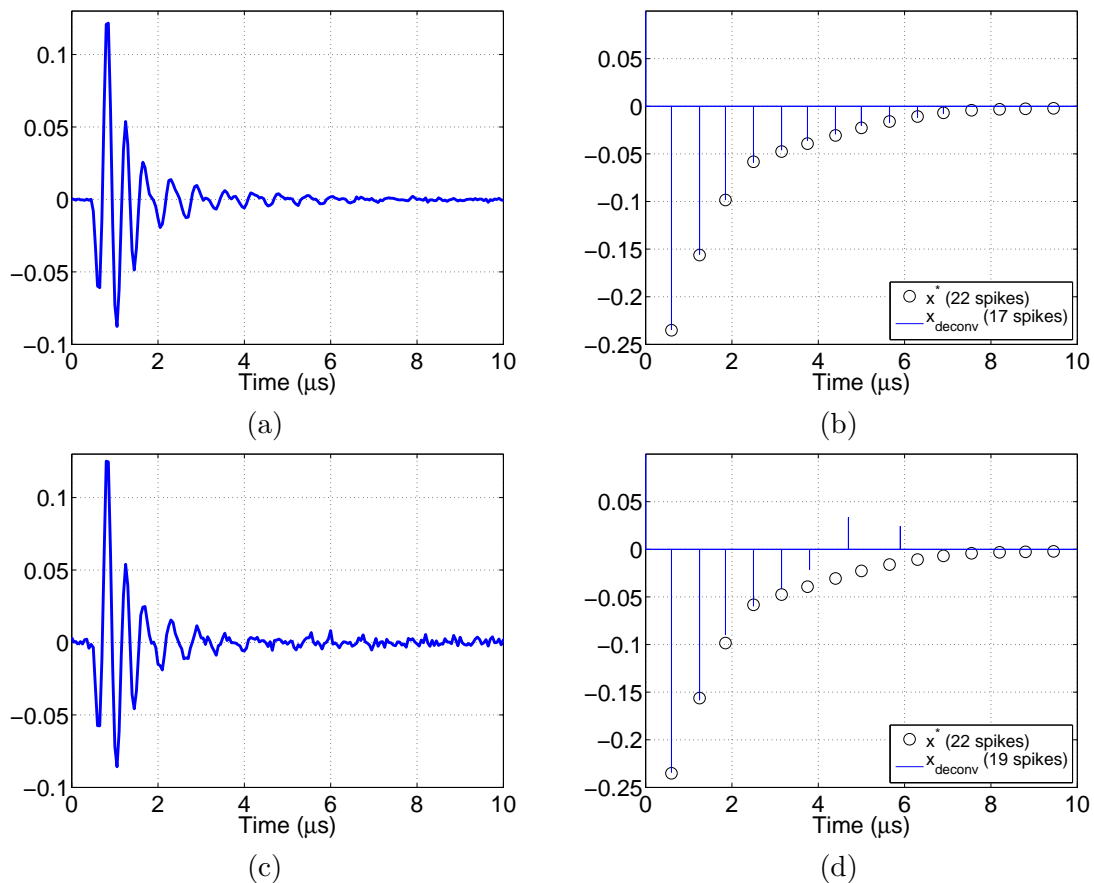

(a)

(b)

(c)

(d)

Figure 4. Deconvolution of acoustic simulated data for two noise levels using the ℓ_0 pseudo-norm deconvolution algorithm. The data signal $\boldsymbol{y}$ (of size 2048) corresponds to the propagation of the wavelet displayed in figure 3 through an homogeneous block of aluminium of depth 2 mm. The incident angle is normal to the block surface. When the noise is low (SNR = 20 dB, simulation (a)), the deconvolution algorithm performs accurate estimation of the spike locations (subfigure (b)). For more noisy data (SNR = 10 dB, simulation (c)), the algorithm still locates the main echoes in a precise manner (subfigure (d)). Further echoes cannot be detected since the useful signal is almost completely drowned in noise.

and 3 seconds, respectively, using a code written in Matlab on a laptop PC, Intel Core 2 Duo CPU 1.4 GHz.

5. Conclusions

Sparse deconvolution methods based on ℓ_0 and ℓ_1 regularisations are efficient tools to process acoustic signals resulting from wave propagation inside a material. The accurate detection and localisation of the time of arrival of the acoustic echoes is critical for detecting and localising the inhomogeneities (change of structures, faults) of the insonified object. For the acoustic data we simulated, we considered a simple, homogeneous material of small thickness. However, the method can also be applied to objects of more complex geometry.

In the near future, our goal will be to consider dispersive and/or attenuative cases. In such cases, the shape of the wavefront evolves as the wave propagates, so that proper adaptation of the wavefront must be performed within the deconvolution method, *i.e.,* deconvolution must become *adaptive*. Olofsson and Stepinski's work [30] is a pioneering contribution in this regard. A source of improvement could be to consider a physical attenuation (or dispersion) model, rather than a signal processing-oriented one. Another one would be to address the problem using the tools of sparse deconvolution introduced in the present paper.

References

[1] Van Trees H L 1968 *Detection, estimation and modulation theory* Part 1 (New York, NY: John Wiley)
[2] Champagnat F, Goussard Y, Gautier S and Idier J 2008 *Bayesian Approach to Inverse Problems* ed Idier J (ISTE Ltd and John Wiley & Sons Inc) pp 117–140
[3] Kormylo J J and Mendel J M 1982 *IEEE Trans. Inf. Theory* **28** 482–488
[4] Mendel J M 1983 *Optimal Seismic Deconvolution* (New York, NY: Academic Press)
[5] Champagnat F, Goussard Y and Idier J 1996 *IEEE Trans. Signal Processing* **44** 2988–2998
[6] Benammar A, Drai R and Guessoum A 2008 *Int. J. Simul. Multidisci. Des. Optim.* **2** 107–111
[7] Taylor H, Banks S and McCoy F 1979 *Geophysics* **44** 39–52
[8] Lorenz D A and Trede D 2009 *Signal Processing with Adaptive Sparse Structured Representations (SPARS workshop)* (Saint-Malo, France) pp 1–6
[9] Soussen C, Idier J, Brie D and Duan J 2010 From Bernoulli-Gaussian deconvolution to sparse signal restoration Tech. rep. Centre de Recherche en Automatique de Nancy
[10] Elad M 2010 *Sparse and Redundant Representations: From Theory to Applications in Signal and Image Processing* (Springer)
[11] Duan J, Soussen C, Brie D and Idier J 2009 *Signal Processing with Adaptive Sparse Structured Representations (SPARS workshop)* (Saint-Malo, France) pp 1–6
[12] Wang Y 2004 *Handbook of Computational Statistics* vol 1 ed Gentle J E, Härdle W and Mori Y (Berlin, Germany: Springer-Verlag) pp 437–466
[13] Idier J (ed) 2008 *Bayesian Approach to Inverse Problems* (ISTE Ltd and John Wiley & Sons Inc)
[14] Pati Y C, Rezaiifar R and Krishnaprasad P S 1993 *Proceedings of 27th Asilomar Conference on Signals, Systems and Computers* vol 1 pp 40–44
[15] Chen S S, Donoho D L and Saunders M A 1998 *SIAM J. Sci. Comput.* **20** 33–61
[16] Goussard Y, Demoment G and Idier J 1990 *Proc. IEEE ICASSP* (Albuquerque, NM) pp 1547–1550
[17] Efron B, Hastie T, Johnstone I and Tibshirani R 2004 *Annals Statist.* **32** 407–451
[18] Donoho D L and Tsaig Y 2008 *IEEE Trans. Inf. Theory* **54** 4789–4812
[19] Stoica P and Selén Y 2004 *IEEE Trans. Signal Processing Mag.* **21** 36–47
[20] Akaike H 1974 *IEEE Trans. Automat. Contr.* **AC-19** 716–723
[21] Rissanen J 1983 *Annals Statist.* **11** 416–431
[22] Hannan E J and Quinn B G 1979 *J. R. Statist. Soc. B* **41** 190–195
[23] Wahba G 1977 *SIAM J. Num. Anal.* **14** 651–667
[24] Golub G H, Heath M and Wahba G 1979 *Technometrics* **21** 215–223
[25] Miller A J 2002 *Subset Selection in Regression* 2nd ed (London, UK: Chapman and Hall)
[26] Hansen M H and Yu B 2001 *J. Acoust. Society America* **96** 746–774
[27] Grünwald P D (ed) 2007 *The Minimum Description Length Principle* MIT Press Books (The MIT Press)
[28] de Ridder F, Pintelon R, Schoukens J and Gillikin D P 2005 *IEEE Trans. Instrum. and Meas.* **54** 144–150
[29] Potel C, de Belleval J F and Gargouri Y 1995 *J. Acoust. Society America* **97** 2815–2825
[30] Olofsson T and Stepinski T 2001 *J. Acoust. Society America* **109** 2831–2839