GEOPHYSICS, VOL. 73, NO. 1 (JANUARY-FEBRUARY 2008); P. R1–R9, 14 FIGS.
10.1190/1.2790584

Stochastic sparse-spike deconvolution

Danilo R. Velis[1]

ABSTRACT

Sparse-spike deconvolution can be viewed as an inverse problem where the locations and amplitudes of a number of spikes (reflectivity) are estimated from noisy data (seismic traces). The main objective is to find the least number of spikes that, when convolved with the available band-limited seismic wavelet estimate, fit the data within a given tolerance error (misfit). The detection of the spikes' time lags is a highly nonlinear optimization problem that can be solved using very fast simulated annealing (SA). Amplitudes are easily estimated using linear least squares at each SA iteration. At this stage, quadratic regularization is used to stabilize the solution, to reduce its nonuniqueness, and to provide meaningful reflectivity sequences, thus avoiding the need to constrain the spikes' time lags and/or amplitudes to force valid solutions. Impedance constraints also can be included at this stage, providing the low frequencies required to recover the acoustic impedance. One advantage of the proposed method over other sparse-spike deconvolution techniques is that the uncertainty of the obtained solutions can be estimated stochastically. Further, errors in the phase of the wavelet estimate are tolerated, for an optimum constant-phase shift is obtained to calibrate the effective wavelet that is present in the data. Results using synthetic data (including simulated data for the Marmousi2 model) and field 3D data show that physically meaningful high-resolution sparse-spike sections can be derived from band-limited noisy data, even when the available wavelet estimate is inaccurate.

INTRODUCTION

The deconvolution of seismic traces is one important digital signal processing method in geophysical exploration. Once the earth model and the forward problem have been set mathematically (for example, using the convolutional model), the deconvolution becomes an inverse problem where the data are the seismic trace s_t and the unknown is the reflectivity r_t, as given by

$$s_t = w_t * r_t + n_t, \quad t = 1,2,\ldots,N, \tag{1}$$

where w_t is the seismic wavelet and n_t is the additive noise. This model assumes that the earth structure can be represented adequately by a set of planar layers of constant impedance. For the purpose of this paper, I also assume the wavelet (or an estimate of the wavelet) is known within a constant-phase shift.

It is clear that because of the bandwidth limitation of the wavelet (and data) and because the data are finite and inaccurate, there exists an infinite number of reflectivities that fit the data equally well. One way to solve this nonuniqueness is to pick, among all possible solutions, those solutions that (1) fit the data, (2) satisfy a given set of constraints (if available), and (3) are more likely to be correct. The last issue represents the key to most inverse problems. It is usually put into practice by assuming some prior information about the type of solution one is interested in. In other words, prior information is used to discard implausible models.

Sparse-spike deconvolution assumes, as a prior, that the reflectivity is a sparse sequence of spikes. The main objective of sparse-spike deconvolution methods is to provide a significant increase in bandwidth content from band-limited seismic observations. This is especially important because the search for subtle hydrocarbons traps has become a major task in today's seismic exploration. Under the assumption of sparseness, the deconvolution problem involves (1) detection of the spikes and (2) estimation of their amplitudes (Velis, 2006). Various methods use different search strategies to locate the spikes and rely on the optimization of different cost functions to satisfy a probabilistic model for the reflectivity (Kormylo and Mendel, 1978; Kaaresen and Taxt, 1998). Other methods proceed to optimize some norm that forces the solution to be sparse (Oldenburg et al., 1983; Debeye and van Riel, 1990; Sacchi et al., 1994; Wang et al., 2006).

In this work, the objective is to find the locations and amplitudes of the least number of spikes that, when convolved with the seismic wavelet, fit the data within a given tolerance error (misfit). Optimum amplitudes are obtained directly using linear least squares (given the

Manuscript received by the Editor 1 March 2007; revised manuscript received 16 August 2007; published online 9 November 2007.
[1]Universidad Nacional de La Plata, Facultad de Ciencias Astronómicas y Geofísicas, and Consejo Nacional de Investigaciones Científicas y Técnicas, La Plata, Argentina. E-mail: velis@fcaglp.unlp.edu.ar.

locations). On the other hand, the detection of the spikes, which represents a highly nonlinear optimization problem, is solved by very fast simulated annealing (SA) (Ingber, 1989). This avoids the suboptimal results of other search strategies.

Vestergaard and Mosegaard (1991) address the problem of deconvolution (through the inversion of poststack data) using sparse prior information and a variant of SA (Nulton and Salamon, 1988; Mosegaard and Vestergaard, 1991). In their method, numerical instabilities in the linear least-squares stage are handled by introducing constraints in the spike locations; for example, two spikes cannot be closer than a predetermined distance. This precludes the use of the procedure in geologic situations such as layer pinch outs and thin beds. Another difficulty mentioned by the authors comes from the fact that the linear stage may lead to solutions with invalid amplitudes. Hence, the use of constrained linear optimization is suggested to alleviate this problem, though it is not clearly explained how these constraints are included from a numerical point of view.

In contrast, the quadratic regularization used in this work yields stable and valid solutions without the need for explicitly constraining either the location or the amplitudes of the reflection coefficients. Furthermore, because the proposed trace-by-trace strategy provides good lateral continuity, the use of additional terms in the cost function is not required in Vestergaard and Mosegaard, 1991 (a weak lateral smoothness constraint is utilized). Once a single trace has been processed, the next trace is deconvolved using the previous solution as the initial model; thus, lateral continuity is obtained automatically.

The proposed methodology allows for any number of impedance constraints to be easily incorporated into the optimization procedure to further reduce the nonuniqueness of the solution and to improve its accuracy and consistency. As a by-product, the acoustic impedance also is recovered if such geologic information is available at some locations and used as constraints for the deconvolution process. Additionally, a strategy for handling wavelet inaccuracies is proposed. Clearly, the success of the deconvolution process is partially affected by the quality of the estimated wavelet used to deconvolve the traces (Levy and Oldenburg, 1982).

Errors in the phase could certainly diminish the results; thus, a strategy for calibrating the wavelet estimate is important. Here, the optimum constant-phase shift is found so that the sparse-spike solution honors the data within a given tolerance for the selected set of parameters (i.e., number of spikes, etc.). As a result, the wavelet estimate is calibrated to carry its phase to that of the effective wavelet governing the data.

One advantage of the proposed deconvolution method is that it is possible to estimate the uncertainty of the inversion by performing several SA runs using different seeds. This procedure allows one to explore stochastically the model space and to obtain a family of acceptable sparse-spike solutions that honor both data and constraints. At the same time, using the average solution as the final solution (instead of an individual realization), the generation of spurious artifacts that may destroy lateral continuity is minimized. The method is illustrated using various synthetically generated (including simulated data for the Marmousi2 model) and field data examples. These examples show that the proposed deconvolution method is able to accurately recover a useful high-resolution sparse reflectivity model from band-limited noisy data, even when the phase of the wavelet estimate contains errors.

One may argue that equation 1 together with the sparse-spike assumption is a rather simple model for an inverse method to rely on and that more complex methods based on 2D or 3D modeling should be used instead. However, practical inversion in the exploration community is still dominated by convolutional models, for the cost of the forward problem step is a real petroleum exploration and production concern. This is especially true for testing stochastic methods in a nonlinear optimization framework, where simple and computationally efficient forward modeling is of paramount importance.

THEORY

A sparse reflectivity sequence can be represented by the following model:

$$r_t = \sum_{j=1}^{M} \alpha_j \delta_{t-\tau_j}, \tag{2}$$

where M, $M \ll N$, is the number of nonzero spikes of amplitudes α_j and time lags τ_j, and δ_t is the delta function. Parameters α_j and τ_j are unknown.

Inserting equation 2 into equation 1 yields

$$s_t = \sum_{j=1}^{M} \alpha_j w_{t-\tau_j+1} + n_t, \quad t = 1,2,\ldots,N. \tag{3}$$

In matrix form, the relationship is

$$\mathbf{A}\boldsymbol{\alpha} + \mathbf{n} = \mathbf{s}, \tag{4}$$

where $\mathbf{A}$ is an $N \times M$ matrix with elements $A_{ij} = w_{i-\tau_j+1}$. This is a set of N nonlinear equations with $2M$ unknowns that can be solved by minimizing the squared error between the actual data (seismic trace) and the calculated data. For this purpose, I define the cost function as

$$J_a = \|\mathbf{A}\boldsymbol{\alpha} - \mathbf{s}\|^2. \tag{5}$$

For simplicity, in this work I assume that the noise (or an estimate of the noise) is known; thus, the expected misfit can be calculated. For normally distributed noise with mean zero and standard deviation σ, the expected misfit is

$$E[J_a] = N\sigma^2. \tag{6}$$

In actual computations, the misfit is calculated using $\tilde{J}_a = \sqrt{J_a/N}$, meaning that the expected misfit is $E[\tilde{J}_a] = \sigma$. Thus, any configuration of M spikes with lags τ_j and amplitudes α_j that satisfies $\tilde{J}_a \leq \sigma$ honors the data and represents a plausible solution. Because the prior is represented by model 2, all solutions that fit the data are guaranteed to be sparse, provided $M \ll N$.

I minimize J_a with respect to the time lags $\hat{\tau}_j$ using very fast SA. Amplitudes $\hat{\alpha}_j$ are estimated by linear least squares (given the lags) at each SA iteration using

$$\frac{\partial J_a}{\partial \hat{\boldsymbol{\alpha}}} = 2\mathbf{A}^T(\mathbf{A}\hat{\boldsymbol{\alpha}} - \mathbf{s}) = 0, \tag{7}$$

which leads to

$$\hat{\boldsymbol{\alpha}} = (\mathbf{A}^T\mathbf{A})^{-1}\mathbf{A}^T\mathbf{s}, \tag{8}$$

where $A_{ij} = w_{i-\hat{\tau}_j+1}$. Because M is small in general, the computational cost associated with equation 8 is low.

Impedance constraints

The importance of using low frequencies to correctly estimate acoustic impedance is well known (Oldenburg et al., 1983). If we assume that a set of K acoustic impedances I_i is available at different time levels T_i, we can use them as constraints in the deconvolution approach to recover the missing low-frequency content. For this purpose, it is useful to use the logarithmic approximation that relates the reflection coefficient with the impedance, which is valid for small $|r_t|$ (less than about 0.3):

$$r_t = \frac{I_t - I_{t-1}}{I_t + I_{t-1}} \simeq \frac{\Delta I_t}{2I_t} = \frac{1}{2}\Delta \log I_t. \quad (9)$$

Then

$$\sum_{j=1}^{T_i} r_j \simeq \frac{1}{2}\log\frac{I_i}{I_0} = \xi_i, \quad i = 1,K, \quad (10)$$

where I_0 is the impedance at the time level T_0 (e.g., the surface) and ξ_i are the impedance constraints at time levels T_i, the values which are assumed to be known. Finally, using the sparse-spike model 2 with estimated lags $\hat{\tau}_j$, I get

$$\sum_{j=1}^{M} \alpha_j u_{T_i - \hat{\tau}_j} = \xi_i, \quad i = 1,K, \quad (11)$$

where u_t is the step function. In matrix form, the equation becomes

$$\mathbf{C}\boldsymbol{\alpha} = \boldsymbol{\xi}, \quad (12)$$

where $\mathbf{C}$ is a $K \times M$ matrix with elements $C_{ij} = u_{T_i - \hat{\tau}_j}$ and $\boldsymbol{\xi}$ is a K-length vector containing the constraints.

I can now write the cost function as

$$J = \|\mathbf{A}\boldsymbol{\alpha} - \mathbf{s}\|^2 + \mu\|\mathbf{C}\boldsymbol{\alpha} - \boldsymbol{\xi}\|^2 = J_a + \mu J_c, \quad (13)$$

where μ is a trade-off parameter that is used to control the weight between the data misfit J_a and the constraint misfit J_c. Thus, for a fixed set of spikes' lags, the least-squares solution is given by

$$\frac{\partial J}{\partial \hat{\boldsymbol{\alpha}}} = 0 \quad \Rightarrow \quad \hat{\boldsymbol{\alpha}} = \mathbf{F}^{-1}(\mathbf{A}^T\mathbf{s} + \mu\mathbf{C}^T\boldsymbol{\xi}), \quad (14)$$

where $\mathbf{F} = \mathbf{A}^T\mathbf{A} + \mu\mathbf{C}^T\mathbf{C}$.

Quadratic regularization

So far, the inverse problem has been regularized by fixing the number of spikes of the sparse-spike model. However, the matrix inversion involved in equation 14 may require additional regularization to stabilize the solution further, especially when M is relatively large. So, the new solution, given the lags, can be written as

$$\hat{\boldsymbol{\alpha}} = (\mathbf{F} + \beta\mathbf{I})^{-1}(\mathbf{A}^T\mathbf{s} + \mu\mathbf{C}^T\boldsymbol{\xi}), \quad (15)$$

where β is the damping factor (also known as prewhitening) and $\mathbf{I}$ is the identity matrix.

It is easy to show that the above solution is obtained after minimizing the new cost function

$$J = J_a + \mu J_c + \beta J_\alpha, \quad (16)$$

where $J_\alpha = \|\boldsymbol{\alpha}\|^2$ is the L_2-norm of the solution. In practice, I set $\beta = \beta_0 f_0$, where $f_0 = \max_i\{F_{ii}\}$ and β_0 is some positive small number. However, for a given M, one would like β_0 to be as large as possible, provided both data and constraints are honored. The selection of M and the trade-off parameters is discussed later.

The algorithm

The algorithm for the sparse-spike deconvolution of the seismic trace with impedance constraints is an iterative process where each SA iteration involves guessing a set of M lags, updating matrices $\mathbf{A}$ and $\mathbf{C}$ (and $\mathbf{F}$), and solving equation 15. Misfit J_a also is computed to check convergence. When the expected misfit and/or a user-defined maximum number of iterations $it_{\max}$ is reached, the SA iteration stops. Step by step, the algorithm is as follows:

1) Fix M. Set the trade-off parameters μ and β.
2) Set the initial set of time lags $\{\hat{\tau}_j^0\}, j = 1,M$.
3) For each SA iteration l, do the following:
   a) Set the new set of time lags $\{\hat{\tau}_j^l\} = \{\hat{\tau}_j^{l-1} + \Delta\hat{\tau}_j^l\}$, $j = 1,M$.
   b) Update $\mathbf{A}$, $\mathbf{C}$, and amplitudes $\hat{\boldsymbol{\alpha}}$ using equation 15.
   c) Evaluate J using equation 16.
   d) According to the Metropolis criterion (as usual in SA), accept or reject the proposed set of time lags.
   e) Also, evaluate $\tilde{J}_a$; if $\tilde{J}_a > \sigma$ and $l < it_{\max}$, go to Step 3a.
4) Convergence (or maximum number of iterations) is achieved.

Steps 1–3 can be repeated using a different number of spikes and/or trade-off parameters. Also, constraints in the location of one or more spikes can be implemented easily. This allows one to introduce some previous knowledge from well-log analysis or data-processing techniques, especially for the most important reflectors. Because each unknown parameter (the sample index along the trace) has an associated search range, it is easy to fix one or more spikes' locations (or narrow their search range) during the SA iteration. This process certainly would reduce the nonuniqueness of the solution and the uncertainty at these locations.

Uncertainty

In spite of M being fixed and β_0 being selected as large as possible, there still may exist many solutions that fit the data equally well (the model space may still be quite large). Logically, the greater the noise level, the greater the uncertainty (variability). In this regard, it is very useful to take advantage of the stochasticity of the SA process and get an estimate of the uncertainty of the obtained solution by performing several SA runs using different seeds. In the next examples, I compute the uncertainty by repeating the SA run using different seeds while keeping M, β_0, and μ fixed. Of course, the data, the constraints, and the noise are the same in all runs.

RESULTS AND DISCUSSION

1D synthetic data

Figure 1 illustrates the proposed method. The first row shows the true reflectivity sequence (composed of 12 spikes), the corresponding impedance profile, the seismic trace, and the wavelet. Gaussian noise with $\sigma = 0.02$ (signal-to-noise ratio of 10) was added to the data (the maximum trace value is about 0.2, so the noise level is about 10% by amplitude). The wavelet is a rotated Ricker wavelet with a central frequency of 30 Hz.

The next row shows the results (red only) after averaging 100 independent SA runs (no impedance constraints). In this case, I selected $M = 15$ and $\beta_0 = 0.2$. All 100 solutions fit the data equally well. In practice, the iteration stopped when the expected misfit was reached. The shaded areas illustrate the uncertainty of the solution (66 out of 100 solutions fall within these areas). The true reflectivity function was recovered quite well, with low uncertainty (except for the smallest spikes). In contrast, the uncertainty and the error on the recovered impedance is large. The third row (red only) shows the average solution and uncertainty when $M = 30$ and $\beta_0 = 0.6$. The differences compared with the $M = 15$ case are not significant, showing that the selection of M was not critical for these data.

For comparison, the fourth row in Figure 1 (red only) shows the nonsparse (band-limited) least-squares solution obtained after solving 15 with $\tau_j = 1, N - L + 1$, where L is the wavelet length (i.e., one spike at every sample). In this case, I selected $\beta_0 = 0.7$, so the final misfit approached the expected misfit. No uncertainties can be computed for the nonsparse least-squares solution, which failed to recover the true reflectivity sequence because its lack of resolution.

Rows two and three (blue only) demonstrate the effect of the impedance constraints on the final solution. For this purpose, I assume that the acoustic impedance is known at $T_1 = 0.2$ s and $T_2 = 0.4$ s. I set $\mu = 10$ to guarantee that the constraints are always honored accurately. For $M = 15$, I set $\beta_0 = 0.02$, and for $M = 30$, $\beta_0 = 0.06$. Finally, the last row (blue only) shows the nonsparse (band-limited) least-squares solution ($\beta_0 = 0.07$). As expected, the accuracy of the results obtained using the proposed deconvolution approach has improved significantly, especially for the recovery of the acoustic impedance.

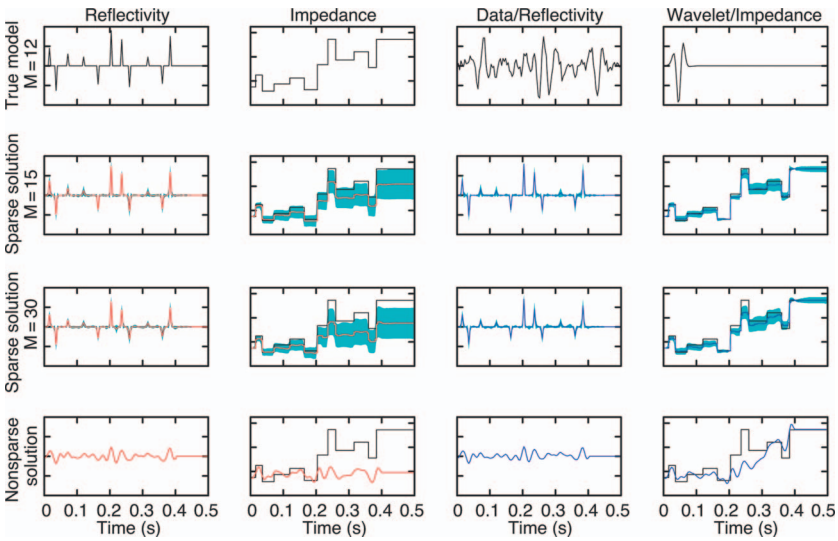

Figure 1. Stochastic sparse-spike deconvolution. First row: true reflectivity, impedance, data (S/N = 10), and wavelet. Second and third rows: recovered reflectivity and impedances after averaging 100 independent runs with $M = 15$ and $M = 30$, respectively (red, unconstrained case; blue, constrained case). Fourth row: nonsparse (band-limited) least squares deconvolution and resulting impedance. True impedance curves (black) are also shown in all impedance panels for comparison. Green areas show uncertainty.

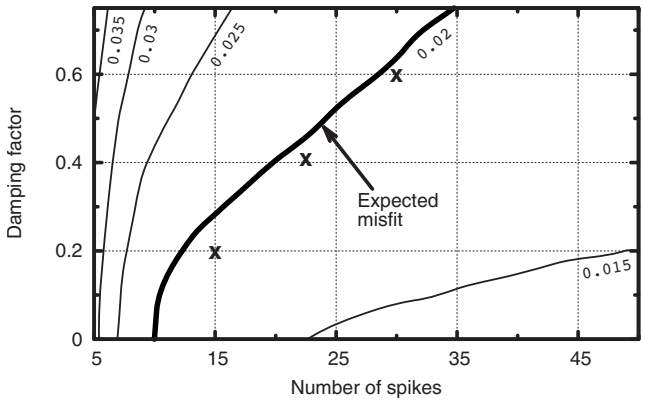

Figure 2. Contour plot of the final misfit (average after 100 realizations) obtained for a range of M and β_0 values in the unconstrained 1D synthetic example. The expected misfit equals 0.02; to the right of this line, data are overfitted; to the left, data are underfitted. The crosses denote preferred pairs of parameters.

Choosing M and the trade-off parameters

The selection of the trade-off parameter associated with the impedance constraints μ usually does not pose any difficulty, for the corresponding term in the cost function μJ_c is not taken into account to check convergence of the algorithm (ultimately, convergence is checked by evaluating the data misfit term). As long as the number of constraints K is significantly smaller than the number of data samples N (and, of course, smaller than M), the minimization of the cost function J would eventually lead to solutions that reproduce the data (at SA convergence), irrespective of the value of μ within a reasonable range (note that J is minimized at every SA iteration). Various tests using synthetic data show that, because J_a and J_c are the same order of magnitude, a value for μ of about 10 is enough to force solutions to fit the impedance constraints very accurately.

The other trade-off parameter, the damping factor β_0, deserves greater care, for the norm of the solution (i.e., the spikes' amplitudes) depends on it. Figure 2 shows a contour plot of the final misfit corresponding to the unconstrained 1D synthetic example of Figure 1 for a wide range of M and β_0 values (after averaging 100 realizations). Clearly, all pairs to the right of the expected misfit contour line lead to solutions that fit the data; those to the left do not. In particular, suppose that β_0 has been fixed and it takes a moderate value. Under these circumstances, if M is chosen with too small a value, data might not be honored. On the other hand, if M is chosen too large, data might be overfitted. To favor simpler models and guarantee sparseness, choose the smallest M so that the observations are reproduced as closely as possible without fitting the noise.

Now suppose that M has been fixed and its value is larger than the actual number of spikes of the model. Under these circumstances, if β_0 is chosen too small, data might be overfitted and the solution might be too oscillatory (unstable). On the other hand, if β_0 is chosen too large, the observations might not be reproduced as expected. A proper selection of β_0 would yield a stable solution that fits the data. In any case, one wants to be as conservative as possible to avoid overfitting and to obtain a family of solutions that honors the observations but is as sparse as possible. The crosses in the figure show

three pairs of parameters, M and β_0, that follow these criteria. The results of two of them were already shown in Figure 1.

Note that these two solutions are almost equally sparse, despite the fact that $M = 15$ in the first case and $M = 30$ in the second case. This is because the selection of the appropriate β_0 (as discussed above) makes the solution rather insensitive to changes in the number of spikes (at least within reasonable ranges of M and β_0). All solutions whose misfit approaches the expected misfit (crosses along the contour line 0.02 in Figure 2) show a similar behavior. In a way, the role of β_0 is to contribute to the sparseness of the solution when M is chosen larger than the actual number of spikes.

In summary, if M is fixed, the trade-off parameter β_0 can be chosen using a simple search so that the resulting misfit after the deconvolution equals its expectation. In practice, this can be done by performing a few tests using different values of β_0. Synthetic tests show that a β_0 in the range 0.0–1.0 usually produces the desired results for moderate M values.

Handling wavelet inaccuracies

Though several methods exist to estimate the seismic wavelet that yields reasonable results even for nonminimum phase sources (e.g., Ulrych et al., 1995; Porsani and Ursin, 2000), often the quality of the derived wavelets is data dependent. Clearly, ignoring wavelet inaccuracies would lead to partially successful deconvolution results. In this regard, the knowledge of the wavelet phase is a critical first step in carrying out the deconvolution (Levy and Oldenburg, 1982). Fortunately, seismic wavelets usually are simpler than we think, and a constant-phase shift may be enough to characterize the main observed structure of a processed wavelet to a very good approximation (Neidell, 1991). This suggests that a constant-phase shift can be included as an additional unknown to calibrate the wavelet estimate and to improve the accuracy of the deconvolution process.

Figure 3 (left column) shows the results of the stochastic sparse-spike deconvolution when the utilized wavelet is in error by a constant-phase shift of -0.5 radians with respect to the original wavelet used to generate the data. Often, in spite of the wavelet being clearly wrong, data can be honored anyway by selecting β_0 small enough and/or M large enough, once again showing the nonuniqueness of the inverse problem. In effect, Figure 4 shows the contour plot of the final misfit corresponding to the previous example [the selected parameter pairs (M,β_0) are marked in the plot]. There exist many sparse solutions that fit the data even with the wrong wavelet. In the previous example, I used $\beta_0 = 0.1$ and 0.5 for $M = 15$ and $M = 30$, respectively, reaching the expected misfit in all cases. Even so, the deconvolution results were partially successful, as shown in Figure 3 (left column).

A comparison of Figure 4 with the contour plot shown in Figure 2 (which utilizes the true wavelet) reveals that, for a given pair (M,β_0), the misfit is larger when the wavelet is affected by errors in the phase. In addition, Figure 5 shows the misfit versus the constant-phase shift ϕ used to calibrate the wavelet for various β_0 and $M = 15$ and $M = 30$. The minima are close to the value that would take the inaccurate wavelet to match the phase of the true wavelet (i.e., $\phi \simeq 0.5$ radians), even when a large value of β_0 precludes reaching the expected misfit for any constant-phase shift.

Tests with phase shifts other than -0.5 radians (not shown) produce the same behavior. As a result of this analysis, one possible strategy for handling inaccuracies in the wavelet is to conduct a search for the constant-phase shift that minimizes the misfit for any reasonable set of parameters (M,β_0). As before, the smallest M values that lead to solutions that fit the data at about the expected misfit (0.02 in this case) are preferred. By inspecting Figure 5, this strategy suggests that a good choice of parameters is $\beta_0 = 0.2$ and $\phi \simeq 0.5$ for the $M = 15$ case and $\beta_0 = 0.6$ and $\phi \simeq 0.5$ for $M = 30$. The results for these two cases are shown in Figure 3 (right column).

2D synthetic data

The algorithm can be applied to 2D (or 3D) data in a trace-by-trace process. Experience shows that the desired coherence in the lateral direction can be obtained without including an additional term in the cost function, thus avoiding the need to deal with another trade-off parameter.

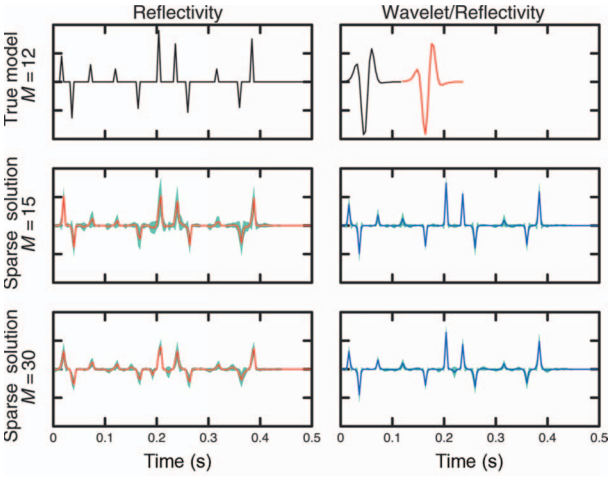

Figure 3. Stochastic sparse-spike deconvolution with wavelet inaccuracies. First row: true reflectivity and wavelets (black, true; red, inaccurate). Second and third rows: recovered reflectivity after averaging 100 independent runs with $M = 15$ and $M = 30$, respectively (red, using the inaccurate wavelet; blue, using the calibrated wavelet). Green areas show uncertainty.

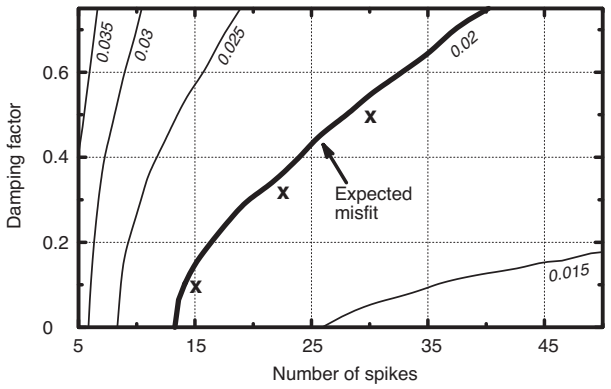

Figure 4. Contour plot of the final misfit (average after 100 realizations) obtained for a range of M and β_0 values in the unconstrained 1D synthetic example. The utilized wavelet is in error by a constant-phase shift of -0.5 radians with respect to the original wavelet.

In the example, the reflectivity section is constructed by extracting a portion of the P-wave velocity Marmousi2 model (Martin et al., 2006), as shown in Figure 6. For the sake of simplicity, I assume constant density throughout the model and constant velocity within each layer, giving rise to the reflectivity section (after depth-to-time conversion) shown in Figure 7b. Then, using the same wavelet of the previous example, I simulate the noise-free zero-offset section shown in Figure 7a. This data set consists of 201 traces with an offset interval of 12.5 m and a time window of 1 s ($\Delta t = 4$ ms).

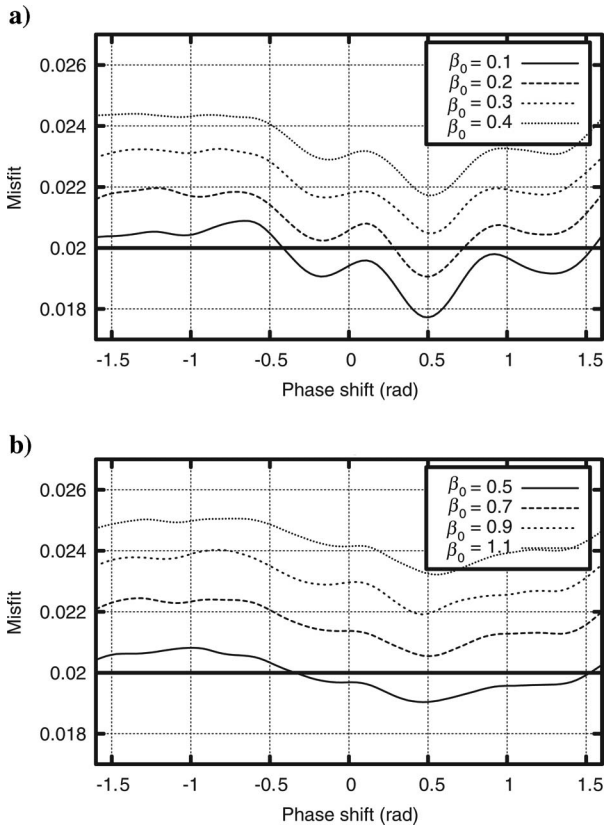

Figure 5. Misfit versus phase shift for various damping factors (β_0) in the unconstrained 1D synthetic example: (a) $M = 15$, (b) $M = 30$.

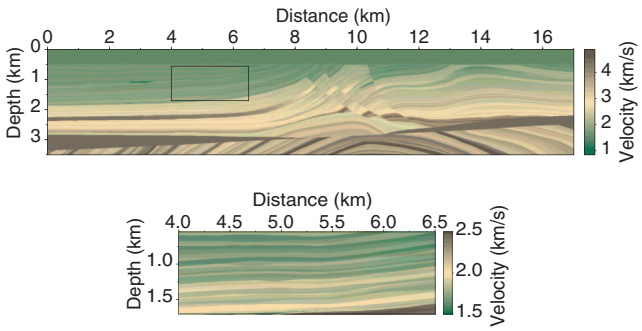

Figure 6. Marmousi2 P-wave velocity model (top) and portion of the original model (bottom) used to generate the reflectivity section shown in Figure 7.

The results of the stochastic sparse-spike deconvolution (noise-free input data) with $M = 45$ and $\beta_0 = 0.25$ are shown in Figure 7c. The estimated reflectivity section is the average section after 100 independent SA runs. Figure 7d illustrates the uncertainty of the derived solution. Despite the fact the M has been selected quite large (the true reflectivity shows no more than 20–25 spikes for any one single trace), the solution does not contain spurious artifacts, and all main reflectors were recovered and resolved very accurately with the appropriate lateral coherence. The five very close reflectors located at about 1.4 s were not resolved so accurately, but in general, the quality of the deconvolution is very good.

Figure 7 (column two) shows the results when the input data are contaminated with 10% additive Gaussian noise. As expected, the uncertainty and variability of the solution increase. Nevertheless, most reflectors are recovered very well.

When the wavelet is in error, the results of the deconvolution deteriorates. Figure 8c shows the results of the sparse-spike deconvolution ($M = 25$, $\beta_0 = 0.1$, S/N = 10) when the wavelet phase is shifted by 0.75 rad. For example, the reflector at about 1.4 s appears as two reflectors instead of one. The wavelet can be calibrated after inspecting Figure 9, where the misfit corresponding to a single trace within the window shown in Figure 8 is plotted for various phase shifts and β_0 ($M = 25$ in all cases). As expected, a phase shift of about -0.75 rad minimizes the misfit for any β_0 (the behavior of the misfit curves for other traces is similar). Finally, the deconvolved section using the calibrated wavelet is shown in Figure 8d, where the true reflectivity is recovered very accurately.

Field data

Figure 10 shows a portion of a seismic 3D migrated data set collected somewhere in the Neuquen basin (Argentina) before and after the sparse-spike deconvolution (for the sake of simplicity, only one inline and crossline are displayed). The utilized seismic wavelet, obtained by standard procedures using the available information provided by a well located within the shown data set ranges, is shown in Figure 11. The selected seismic volume consists of 33×33 traces and a time window of 1 s that goes from 0.8 to 1.8 s ($\Delta t = 2$ ms). For this example, I set $M = 25$ and $\beta_0 = 0.05$.

Figure 12 shows a comparison between the amplitude spectra before and after the sparse-spike deconvolution. The spectra were calculated after averaging all of the individual spectra of the analyzed volume. In general, an increase in the frequency content is observed (both low and high frequencies), leading to a significant improvement on the vertical resolution.

The dominant frequency, calculated by averaging the frequencies corresponding to both spectrum end points at -10 dB, goes from about 23 Hz in the original data set to about 45 Hz after the deconvolution. In addition, the overall shape of the spectrum within the seismic main frequency band is preserved and enhanced. It is important to note that the relative amplitudes and lateral coherence of the reflectors also are well preserved throughout the deconvolved traces. One feature of interest is the anomalous region marked in Figure 10 with a green oval, which is easier to delineate after the sparse-spike deconvolution.

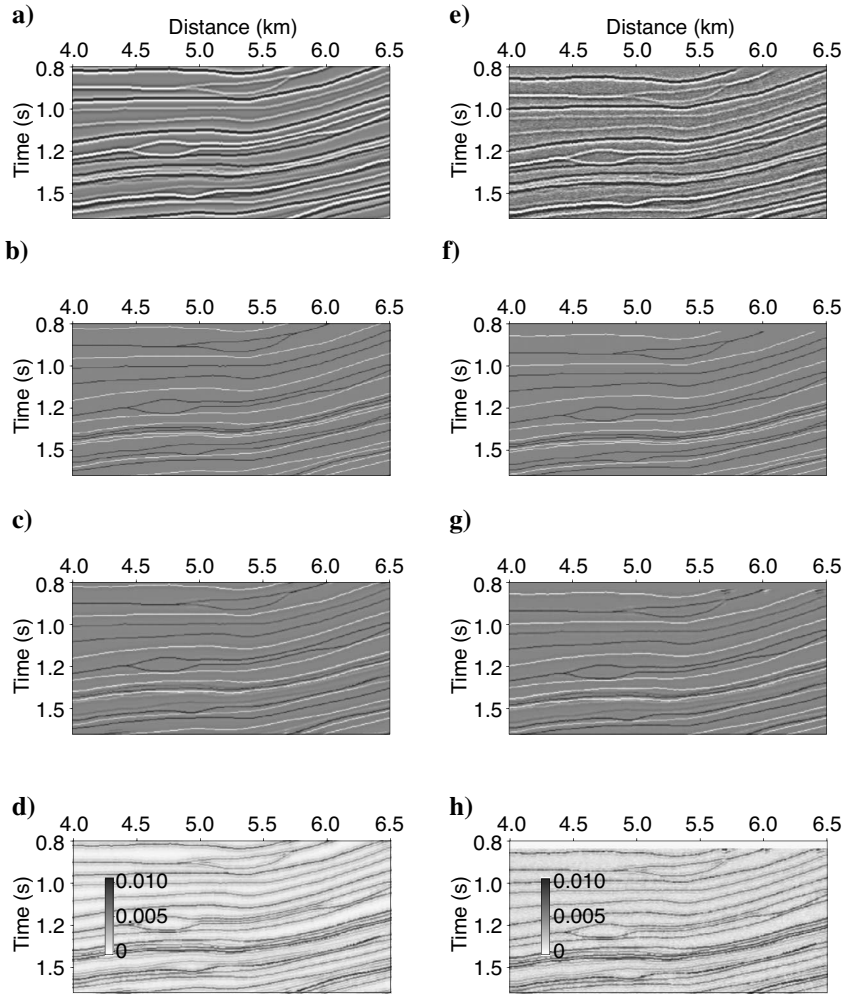

Figure 7. Stochastic sparse-spike deconvolution on the Marmousi2 model. Column one, noise-free case; column two, noisy case (SNR = 10). (a and e) Input data. (b and f) True reflectivity section. (c and g) Estimated reflectivity section (average after 100 runs). (d and h) Standard deviation of the reflectivity in (c and g).

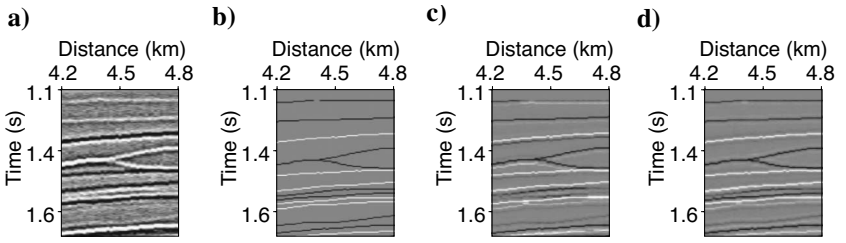

Figure 8. Stochastic sparse-spike deconvolution on the Marmousi2 model. (a) Input data (S/N = 10). (b) True reflectivity. (c) Estimated reflectivity using the wrong wavelet. (d) Estimated reflectivity using the calibrated wavelet.

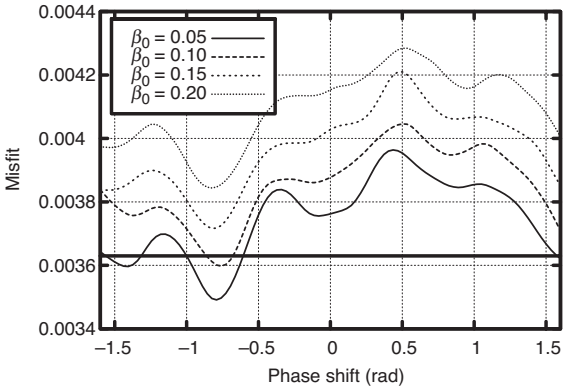

Figure 9. Misfit versus phase shift for various damping factors (β_0) in the Marmousi2 model using $M = 25$ in all cases.

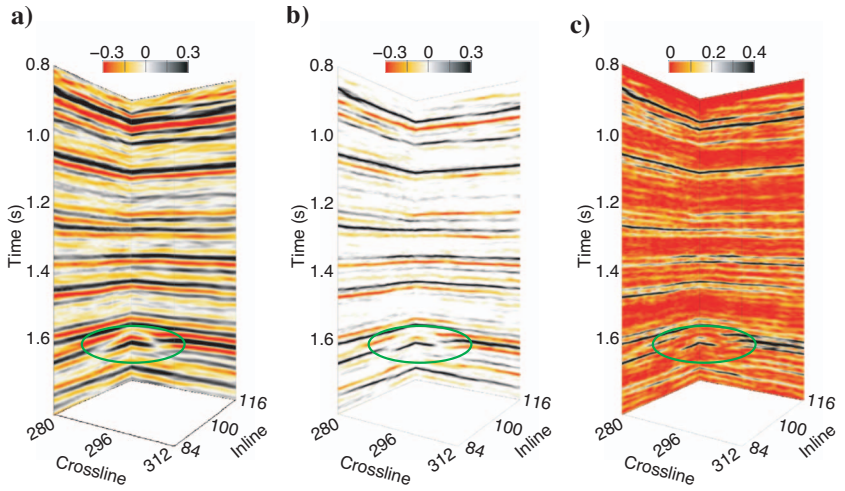

Figure 10. Stochastic sparse-spike deconvolution of a field 3D data set. (a) Input data. (b) Estimated reflectivity using $M = 25$ (average after 100 realizations). (c) Standard deviation of the data in (b). The green oval shows one feature of interest.

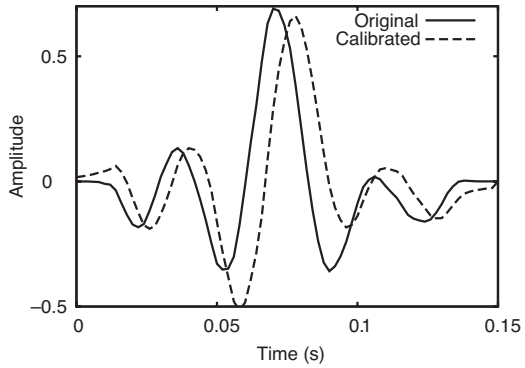

Figure 11. Original zero-phase (solid) and calibrated (dashed) wavelets used to deconvolve the field data set shown in Figures 10a and 14a, respectively.

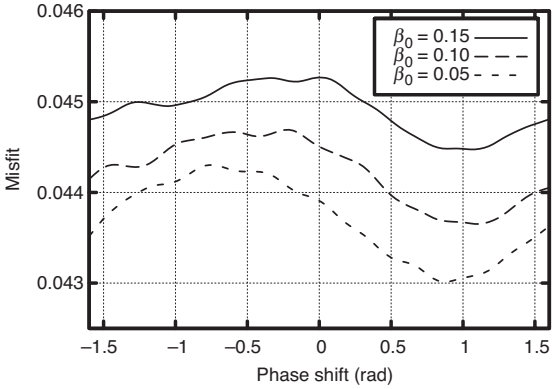

Figure 13. Average misfit versus phase shift for various damping factors (β_0) corresponding to the selected window of the field data set (Figure 14a). $M = 10$ in all cases.

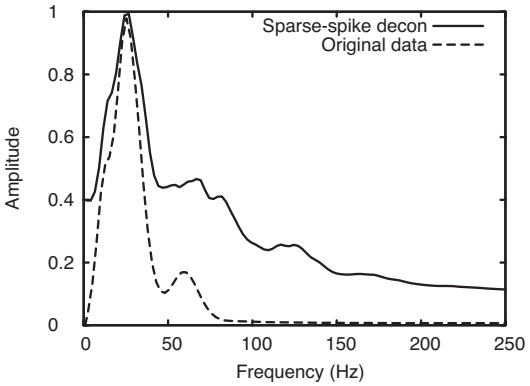

Figure 12. Comparison of the amplitude spectra before and after the application of the sparse-spike deconvolution.

The quality of the deconvolved section can be improved by calibrating the wavelet. I performed this calibration for the 33 traces contained in the inline shown in Figure 10a and a time window in the range of 1.5 to 1.8 s (Figure 14a). Figure 13 shows the average misfit curves for various β_0 ($M = 10$) corresponding to the selected window. The plot reveals that a constant phase shift of about 1.0 rad yields smaller misfit values. Thus, the original zero-phase wavelet was calibrated using this value (Figure 11) and then input to the sparse-spike deconvolution procedure. The results are shown in Figures 14b and c. Note how the lateral continuity of the marked reflectors is improved.

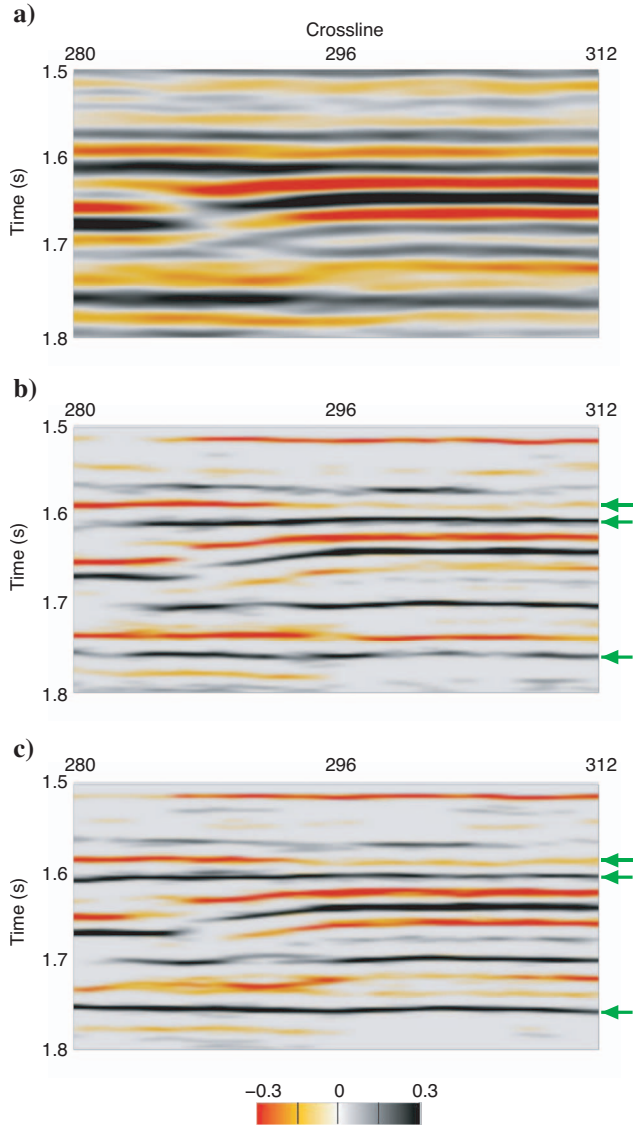

Figure 14. Stochastic sparse-spike deconvolution of the selected window of the field data set ($M = 10, \beta_0 = 0.1$). (a) Input data. (b) Estimated reflectivity using the original zero-phase wavelet (average after 100 realizations). (c) Estimated reflectivity using the calibrated wavelet (average after 100 realizations). Green arrows show reflectors with improved lateral continuity.

CONCLUSIONS

Sparse-spike deconvolution is viewed as a detection and estimation problem. The detection of the time lags of a given number of spikes leads to a highly nonlinear optimization procedure that is solved using very fast SA. The estimation of the amplitudes is done after solving a linear least-squares inverse problem that includes quadratic regularization. Impedance constraints also can be included at this stage.

Quadratic regularization provides the remedy for numerical instabilities and guarantees that the derived reflectivity sequences are meaningful from a physical point of view. In the tested examples, the largest damping factor still leading to solutions that fit the data (and constraints) within a given tolerance error (misfit) is selected after a few trials. This strategy makes the selection of the number of spikes not critical and the variability of the derived solutions smaller. Further, wavelet phase inaccuracies are tolerated, and the corresponding optimum constant-phase shift is derived to calibrate the available wavelet estimate and to improve the quality of the sparse-spike deconvolution.

One advantage of the proposed method is that the uncertainty of the obtained reflectivities can be estimated by running several SA runs using different seeds. The acoustic impedance also is estimated as a by-product of the deconvolution procedure, provided some constraints are included at the linear least-squares stage. As expected, the accuracy of these estimates increases when some impedance constraints at different time levels are introduced. Good lateral continuity is achieved, and no spurious artifacts are generated when processing 2D and 3D data.

ACKNOWLEDGMENTS

I am grateful to Juan Soldo and Repsol YPF (Argentina) for giving me access to the field data set. Financial support for this work was provided by Universidad Nacional de La Plata and Consejo Nacional de Investigaciones Científicas y Técnicas (CONICET), Argentina.

REFERENCES

Debeye, H. W. J., and P. van Riel, 1990, L_p-norm deconvolution: Geophysical Prospecting, **38**, 381–403.
Ingber, L., 1989, Very fast simulated re-annealing: Journal of Mathematical Computation and Modelling, **12**, 967–973.
Kaaresen, K. F., and T. Taxt, 1998, Multichannel blind deconvolution of seismic signals: Geophysics, **63**, 2093–2107.
Kormylo, J., and J. Mendel, 1978, On maximum-likelihood detection and estimation of reflection coefficients: 48th Annual International Meeting, SEG, Expanded Abstracts, 45–46.
Levy, S., and D. Oldenburg, 1982, The deconvolution of phase-shifted wavelets: Geophysics, **47**, 1285–1294.
Martin, G. S., R. Wiley, and K. J. Marfurt, 2006, Marmousi2: An elastic upgrade for Marmousi: The Leading Edge, **25**, 156–166.
Mosegaard, K., and P. Vestergaard, 1991, A simulated annealing approach to seismic model optimization with sparse prior information: Geophysical Prospecting, **39**, 599–612.
Neidell, N., 1991, Could the processed seismic wavelet be simpler than we think?: Geophysics, **56**, 681–690.
Nulton, J., and P. Salamon, 1988, Statistical mechanics of combinatorial optimization: Physical Reviews A, **37**, 1351–1356.
Oldenburg, D. W., T. Scheuer, and S. Levy, 1983, Recovery of the acoustic impedance from reflection seismograms: Geophysics, **48**, 1318–1337.
Porsani, M. J., and B. Ursin, 2000, Mixed-phase deconvolution and wavelet estimation: The Leading Edge, **19**, 76–79.
Sacchi, M. D., D. R. Velis, and A. H. Comínguez, 1994, Minimum entropy deconvolution with frequency domain constraints: Geophysics, **6**, 938–945.
Ulrych, T. J., D. R. Velis, and M. D. Sacchi, 1995, Wavelet estimation revisited: The Leading Edge, **10**, 1139–1143.
Velis, D. R., 2006, Parametric sparse-spike deconvolution and the recovery of the acoustic impedance: 76th Annual International Meeting, SEG, Expanded Abstracts, 2141–2144.
Vestergaard, P., and K. Mosegaard, 1991, Inversion of post-stack seismic data using simulated annealing: Geophysical Prospecting, **39**, 613–624.
Wang, J., X. Wang, and M. Perz, 2006, Structure preserving regularization for sparse deconvolution: 76th Annual International Meeting, SEG, Expanded Abstracts, 2072–2076.