

Seismic sparse-layer reflectivity inversion using basis pursuit decomposition

Rui Zhang¹ and John Castagna²

ABSTRACT

A basis pursuit inversion of seismic reflection data for reflection coefficients is introduced as an alternative method of incorporating a priori information in the seismic inversion process. The inversion is accomplished by building a dictionary of functions representing reflectivity patterns and constituting the seismic trace as a superposition of these patterns. Basis pursuit decomposition finds a sparse number of reflection responses that sum to form the seismic trace. When the dictionary of functions is chosen to be a wedge-model of reflection coefficient pairs convolved with the seismic wavelet, the resulting reflectivity inversion is a sparse-layer inversion, rather than a sparse-spike inversion. Synthetic tests suggest that a sparse-layer inversion using basis pursuit can better resolve thin beds than a comparable sparse-spike inversion. Application to field data indicates that sparse-layer inversion results in the potentially improved detectability and resolution of some thin layers and reveals apparent stratigraphic features that are not readily seen on conventional seismic sections.

INTRODUCTION

In conventional seismic deconvolution, the seismogram is convolved with a wavelet inverse filter to yield band-limited reflectivity. The output reflectivity is band limited to the original frequency band of the data so as to avoid blowing up noise at frequencies with little or no signal. It has long been established (e.g., [Riel and Berkhout, 1985](#)) that sparse seismic inversion methods can produce output reflectivity solutions that contain frequencies that are not contained in the original signal without necessarily magnifying noise at those frequencies. It is well known (e.g., [Tarantola, 2004](#))

that applying valid constraints in seismic inversion can stably increase the bandwidth of the solution.

However, incorporation of the a priori information in the reflectivity inversion of seismic traces can be problematic. A common way of incorporating prior knowledge is to build a starting model biased by that information and to let the inversion process perturb the initial starting model and converge to a solution (e.g., [Cooke and Schneider, 1983](#)). The individual layers represented in the starting model can have hard or soft constraints assigned. This kind of method can work very well when the starting model is close to the correct solution. Typically, the starting model is obtained by spatially interpolating well logs along selected horizons. Unfortunately, these horizons must be picked on the original seismic data. If waveform interference patterns change laterally, horizon picks on a constant portion of a waveform (typically chosen to be peaks, troughs, or zero crossings) can be in error, resulting in an incorrect starting model and a potentially erroneous inversion. Similarly, if velocities and/or impedances for the inversion interval change laterally in a manner different from that resulting from the interpolation procedure, interpolated well logs may again be significantly in error, and the inverse process may converge to the wrong minimum. These problems may be ameliorated with a Monte Carlo approach, but such an approach cannot correct the fundamental nonuniqueness of the process that may cause minimums other than the correct one to have similar errors. A means of biasing the results toward expected reflectivity patterns is needed without relying on possibly erroneous manual interpretations or spatial interpolations.

[Nguyen and Castagna \(2010\)](#) used matching pursuit decomposition (MPD) to decompose a seismic trace into a superposition of reflectivity patterns observed in and derived from existing well control. Matching pursuit decomposition (1) correlates a wavelet dictionary against a seismogram and finds the location, scale (i.e., center frequency), and amplitude of the best-fit wavelet, (2) subtracts the best-fit wavelet and records its characteristics in a table, and (3) repeats the processes on the residual trace until the residual energy falls below a selected threshold. For spectral decomposition,

Manuscript received by the Editor 14 March 2011; revised manuscript received 24 July 2011; published online 22 December 2011.

¹University of Texas, Jackson School of Geosciences, Institute for Geophysics, Austin, Texas, USA. E-mail: rzhang@mail.utexas.edu.

²University of Houston, Department of Earth and Atmospheric Sciences, Houston, Texas, USA. E-mail: jpcastagna@uh.edu.

© 2011 Society of Exploration Geophysicists. All rights reserved.

the spectra of the best-fit wavelets from the table are summed as a function of time to form the time-frequency analysis. For seismic inversion, the wavelet dictionary consists of seismic reflection patterns derived from well logs, which are matched to the seismic trace. In effect, pattern recognition is used to recognize seismic patterns derived from well logs, resulting in what is equivalent to a data adaptive starting model. This ensures that the starting model is not misaligned with the seismic data to be inverted, as may occur when well logs are spatially interpolated. The matching pursuit method has some limitations, especially when the dictionary elements are not orthogonal. Nguyen and Castagna (2010) consequently found lateral instability in the raw MPD solution and had to employ reprojection to obtain laterally stable results. Wang (2007, 2010) developed a multichannel MPD (MCMP) spectral decomposition assuming some degree of lateral coherence to improve the uniqueness and spatial continuity of matching pursuit spectral decomposition. As applied to the seismic inversion problem, such an approach has potential limitations as it (1) could result in some loss of spatial resolution when the geology is, in fact, not continuous and (2) can still exhibit hard lateral jumps when the path-dependent MPD algorithm switches from one initial match to another. For seismic inversion purposes, a more laterally stable approach that operates on a single trace at a time, without the need for a posteriori selection of laterally consistent possible solutions, would be advantageous.

Basis pursuit decomposition (BPD) has many advantages over MPD (Chen et al., 2001); it handles interferences between dictionary elements better, it is computationally more efficient, and, by introducing a sparsity norm and regularization parameter into the objective functions, it can exhibit good lateral stability even when dictionary elements are not orthogonal.

In this study, we investigate the use of BPD to perform seismic inversion. We use the algorithm of Chen et al. (2001) to decompose the seismic trace with a nonorthogonal wavelet dictionary consisting of even and odd thin layer seismic responses. The output reflectivity series is then formed by summing reflection coefficient pairs that are shifted and multiplied by scalars with translations and coefficients output by the basis pursuit algorithm. The basis pursuit algorithm is an L1 norm optimization that was originally developed as a compressive sensing technique. The solution consists of three parts: (1) linear programming transforms the L1 optimization problem into a constrained least square problem, (2) duality theory sets up an array containing primal, dual, and gap equations to be solved, and (3) a primal-dual log-barrier algorithm implements a Gauss-Newton step workflow to solve the equations array.

Much of the advantage of BPD over MPD is related to the fact that BPD finds a single global solution, whereas MPD is a path-dependent process. As MPD iteratively subtracts matched wavelets from the seismogram, the order in which wavelets are subtracted may vary greatly between similar seismic traces. A change in the order in which the wavelets are subtracted can result in an entirely different solution. Slight variations in adjacent input traces (as could result from noise) could cause geologically unreasonable jumps from one solution to another. Furthermore, for interfering wavelets, the location and center frequency of the best wavelet match found by MPD may not correspond to the time location or center frequency of the reflecting seismic wavelets. If a wavelet of the wrong frequency is subtracted, nearby wavelets found by MPD must compensate for this error, often by being wrong in

the opposite direction, and can result in the addition of small satellite wavelets to the solution that are needed to minimize the resulting error. In data compression applications, this results in less than optimal compression but is not necessarily a serious problem. In spectral decomposition using MPD, however, this can cause time asymmetry in the time-frequency decomposition of thin layers, even when the input is time-symmetrical. This results in serious lateral instability in spectral decomposition and inversion using trace-by-trace MDP (Wang, 2010; Nguyen and Castagna, 2010). To avoid these problems, we investigate the use of basis pursuit inversion (BPI) to invert a seismic trace for a reflectivity series which can be integrated to give the band-limited seismic impedance model.

As a first step, we study the use of a very basic reflection pattern dictionary consisting of a wedge-model of reflection coefficient pairs convolved with the seismic wavelet. As any pair of reflection coefficients can be represented as the sum of odd and even reflection coefficient pairs (e.g., Puryear and Castagna, 2008), there will be two dictionary elements (odd and even) for each thickness represented in the wedge-model and the number of elements in the dictionary will be twice the number of thicknesses. The thickness range built into the dictionary then becomes a constraint on the possible outcomes. The direct use of more complex well-log derived reflectivity patterns in the BPI dictionary is left as an objective of future research.

The use of BPI with a wedge dictionary is essentially a sparse-layer inversion. It can be argued that such an inversion may resolve thin beds differently than a sparse-spike type of inversion (e.g., Oldenburg et al., 1983), which imposes some specified degree of sparsity and, thus, indirectly places some limitation on the spacing of reflections. In this paper we will compare BPI sparse-layer inversion with conventional sparse-spike inversion on synthetic and real data to assess the relative ability of BPI to resolve thin beds and reveal fine stratigraphic features.

BASIC THEORY

For simplicity, the forward model is assumed to be a simple convolution of a stationary seismic wavelet and the reflectivity. The seismic trace $s(t)$ is thus given by

$$s(t) = w(t) * r(t) + n(t), \quad (1)$$

where $w(t)$ is the seismic wavelet, $r(t)$ is the reflectivity series, and $n(t)$ is the noise. This model assumes that the earth structure can be represented adequately by a series of planar horizontal layers of constant impedance with reflections generated at the boundaries between adjacent layers. To further simplify the inversion, we assume that the wavelet is known.

Because seismic wavelets are band limited and seismic data are finite and inaccurate, there exists an infinite number of reflectivity series that can fit the data equally well. Thus, the inversion of $s(t)$ for $r(t)$ is nonunique. To find a "best" solution among all possibilities, additional information is required. This best solution should fit the data within some specified tolerance and best satisfy a given set of constraints.

The need to fit the data leads to a least-squares solution. The application of constraints is a key aspect of most inverse problems. Such constraints usually assume some prior information about the type of solution one desires. Prior information is used to discard or weight against implausible possible solutions.

The inverse problem of determining the reflectivity series from a seismic trace and a known wavelet involves two tasks: (1) location of the reflection coefficients and (2) estimation of their amplitudes. 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, 1980; Kaarensen and Tøxt, 1998). Other methods proceed to optimize some norm that forces the results to be sparse (Oldenburg et al., 1982; Riel and Berkhou, 1985; Debeye and van Riel, 1990; Sacchi et al., 1994; Hu et al., 2008). Mosegaard and Vestergaard, 1991 address the problem using sparse prior information. In their methods, 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.

In exploration geophysics practice, conventional sparse-spike inversion (SSI) is commonly used to produce sparse reflectivity series without the need for a starting model. Here, the term “sparse” is used to describe a time series that has a relatively small number of nonzero values. All else being constant, the greater the sparseness, the fewer the number of layers in the inversion result. This limits the spacing between reflection coefficients, thereby potentially discriminating against solutions with thin layers. This possible trade-off between sparseness and resolution can potentially restrict the ability of SSI to resolve thin layers.

Partyka (2005), Portniaguine and Castagna (2005), and Puryear and Castagna (2008) discuss thin-bed inversion in the frequency domain. Partyka coined the term “spectral inversion” for this procedure. Chopra et al. (2006, 2009) show many examples of the use of spectral inversion. Puryear and Castagna (2008) concentrate on the application of spectral inversion to the problem of thin layer thickness determination. Spectral inversion utilizes spectral decomposition to localize time-varying spectral interference patterns created by a limited number of superposed layer responses and to invert these local frequency spectra for the layer thicknesses and reflection coefficients. The application of such a sparse-layer inversion in the frequency domain is complicated by windowing effects, and it could be advantageous to perform equivalent operations directly in the time domain. This can be accomplished using basis pursuit decomposition into individual layer responses. However, frequency-domain implementation can have greater flexibility and computational efficiency than our time-domain implementation, in that different frequency bands (or other segmentations of the spectrum) can be weighted differently with ease. Our intention here is not to show that basis pursuit decomposition is superior or comparable to any particular spectral inversion algorithm (we make no such claim in that regard) but, rather, to present it as an alternative approach that, with further development (e.g., following Nguyen and Castagna, 2010), could more readily incorporate well reflectivity patterns. We compare the basis pursuit inversion method in detail to conventional sparse-spike inversion to show the potential added value in a sparse-layer inversion, irrespective of exactly how it is implemented.

BASIS PURSUIT INVERSION (BPI)

The subsurface is assumed to consist of horizontal isotropic homogeneous layers. Each trace of the poststack seismic image is considered to be the convolution of the seismic wavelet with a reflectivity series. To allow the seismic wavelet to vary temporally,

the convolutional model (equation 1) can also be written for seismic reflection data as

$$\mathbf{s} = \mathbf{W}\mathbf{r} + \mathbf{n}, \quad (2)$$

where \mathbf{s} is a column vector representing the seismogram, \mathbf{r} is the reflectivity series column vector, \mathbf{W} represents the diagonal wavelet kernel matrix, and \mathbf{n} is a noise vector. This convolutional operation has the well-known form:

$$\mathbf{d} = \mathbf{G}\mathbf{m} + \mathbf{n} \quad (3)$$

where \mathbf{d} is the data vector, \mathbf{m} is the model parameters, \mathbf{G} is the kernel, and \mathbf{n} is the noise.

Basis pursuit (BP) solves for parameters in equation 3 by simultaneously minimizing both the L_2 norm of error term and the L_1 norm of the solution:

$$\min[\|\mathbf{d} - \mathbf{G}\mathbf{m}\|_2 + \lambda\|\mathbf{m}\|_1]. \quad (4)$$

The basis pursuit algorithm used in this paper is described in detail by Chen et al. (2001). Here we describe how basis pursuit is used to perform seismic inversion.

Any reflectivity series can be decomposed into a summation of impulse pairs (Bork and Wood, 2001). BPI utilizes dipole decomposition to represent the reflectivity series as a sum of even and odd impulse pairs multiplied by scalars.

The top and base reflectors of a layer can be represented as two impulse functions $c\delta(t)$ and $d\delta(t + n\Delta t)$, where $n\Delta t$ is time-thickness of thin-bed, Δt is sample rate, and c and d are the two reflection coefficients. Dipole decomposition is used to decompose each reflector pair into one even pair r_e and one odd pair r_o with corresponding coefficients a and b , varying from -1 to $+1$ times a scale factor, which can be expressed as equations 5 and 6, also in Figure 1a:

$$r_e = \delta(t) + \delta(t + \mathbf{n}\Delta t); r_o = \delta(t) - \delta(t + \mathbf{n}\Delta t); \quad (5)$$

$$c\delta(t) + d\delta(t + \mathbf{n}\Delta t) = ar_e + br_o. \quad (6)$$

The layer thickness is typically unknown. To include all possible bed thicknesses, \mathbf{n} varies from zero to the number \mathbf{N} with $\mathbf{N}\Delta t$ representing the maximum layer time-thickness. This comprises the wedge dictionary. Thicknesses greater than the maximum thickness are readily accommodated by including a single impulse in the dictionary (this corresponds to a zero thickness even impulse pair).

The wedge models, including the odd and even wedges, are a collection of dipole reflectors with an increasing time separation (Figure 1b – Figure 1e). Figure 1b and 1c displays odd and even wedge reflectivity models, respectively. Figure 1d and 1e shows corresponding seismic responses after convolution with a given wavelet. Interference from the top and base of the wedge models, which degrades the seismic resolution, becomes useful information for BPI.

Because the sample rate is Δt , each trace of even wedge reflectivity consists of a pair of equal impulse functions (spikes) with interval $n\Delta t$. The reflector kernel matrix for the reflectivity pair is constructed by shifting the reflectivity pair along the time axis

with $m\Delta t$, where m ranges from one to the number of samples in the seismic trace. So each even wedge reflectivity can be written as

$$r_e(t, m, n, \Delta t) = \delta(t - m\Delta t) + \delta(t - m\Delta t + n\Delta t). \quad (7)$$

The odd wedge reflectivity has the same pattern as even wedge with the exception of polarized dipoles, written as

$$r_o(t, m, n, \Delta t) = \delta(t - m\Delta t) - \delta(t - m\Delta t + n\Delta t). \quad (8)$$

Any reflectivity series can be considered a summation of even and odd wedge reflectivity patterns, as shown in equation 9:

$$r(t) = \sum_{n=1}^N \sum_{m=1}^M (a_{n,m} * r_e(t, m, n, \Delta t) + b_{n,m} * r_o(t, m, n, \Delta t)). \quad (9)$$

When convolved with the wavelet, w , the left-hand side of equation 9 becomes the seismic trace $s(t)$, and the right-hand side of equation 9 becomes the summation of the even and odd wedge seismic responses, as shown in equation 10:

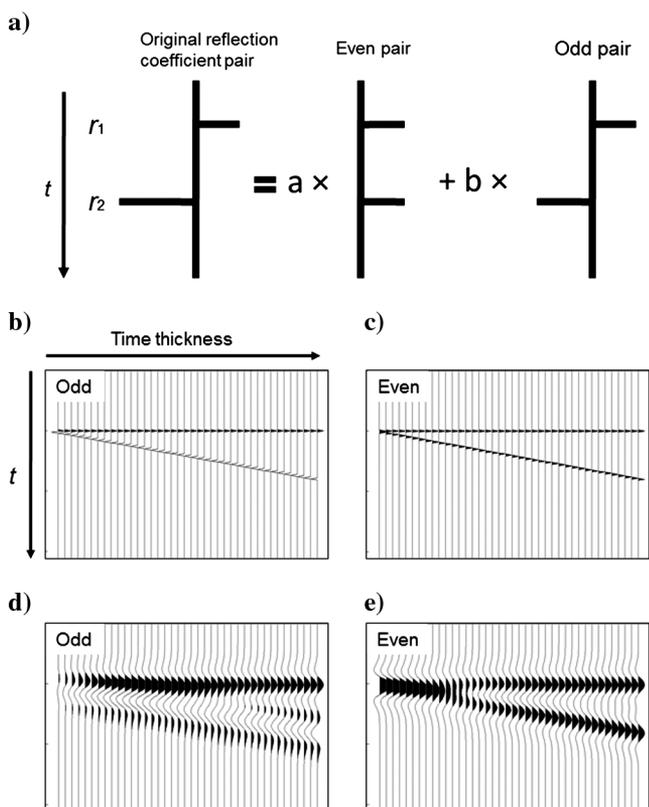


Figure 1. (a) Any arbitrary pair of reflection coefficients r_1 and r_2 can be represented as the sum of even and odd components. The even pair has the same magnitude and sign, and the odd pair has the same magnitude and opposite sign; (b) and (c) are odd and even wedge reflectivity pairs; (d) and (e) show the BPI dictionary elements which are the wedge-model seismic responses using 30 Hz Ricker wavelet.

$$s(t) = \sum_{n=1}^N \sum_{m=1}^M (a_{n,m} * wr_e(t, m, n, \Delta t) + b_{n,m} * wr_o(t, m, n, \Delta t)), \quad (10)$$

where wr_e and wr_o are seismic responses of wedge reflectivity pairs. Any seismic trace can be considered the summation of wedge seismic dictionary elements. Assuming the seismic trace is properly scaled to reflection coefficient magnitudes, the decomposition of reflectivity into wedge reflector pairs and of the seismic trace into wedge seismic responses shares exactly the same scalar coefficients $a_{n,m}$ and $b_{n,m}$.

BPI can be used to calculate coefficients $a_{n,m}$ and $b_{n,m}$ by solving equation 10. As equation 9 and 10 share the same coefficients, the inverted series can be calculated by summing up the wedge reflectivity models with coefficients $a_{n,m}$ and $b_{n,m}$, as illustrated in equation 9. Because the reflection coefficients and dipole reflector pairs are defined with values ranging from -1 to $+1$, the coefficients $a_{n,m}$ and $b_{n,m}$ are of the same magnitude as the reflection coefficients if the waveform amplitudes are properly scaled to reflection coefficients. In practice, the actual scaling of the waveform amplitudes is not a trivial process, and time and spatially varying scalars are often required. We attempt to scale the data to synthetic seismograms at well locations and balance the amplitudes between wells by assuming that particular geological intervals have, on the average, constant or slowly varying reflection coefficients. If we are confident in our amplitude scaling, we can further constrain the solution based on petrophysically determined distributions of reflection coefficients. In this paper we did not apply scaling, beyond normalizing the wavelet by its rms amplitude, nor did we apply petrophysical constraints, to keep the comparison between our method and sparse-spike inversion as simple as possible.

BPI implementation relies on the correct wavelet, which is usually derived from seismic data. Though several methods exist to estimate the seismic wavelet that yields reasonable results (e.g., Ulrych et al., 1995; Porsani and Ursin, 2000), the quality of the derived wavelets is data dependent. A detailed, sensitive test of the wavelet is shown in Zhang (2010). Generally, impedance inversion results are no better than the estimate of the wavelet, and the greater the resolution of the output, the more crucial it is to start with a good estimate of the amplitude and phase spectrum of the wavelet. This is true of all inversion methods; however, sparse inversion methods, in particular, are highly sensitive to the low- and high-frequency wavelet roll-offs (spectral slopes). If the roll-off is too sharp, noise at those frequencies will be amplified. Conversely, if the wavelet amplitude spectrum is too broad (as may be caused by incomplete cancellation of the reflectivity spectrum in wavelet extraction) the inverted reflectivity spectrum can exhibit spectral notches at frequencies where energy is falsely attributed to the wavelet. If the wavelet phase is wrong, this will impact not only the phase of the inversion result but also affect the statistics of the output reflectivity.

SYNTHETIC EXAMPLES

The regularization parameter λ in equation 4 balances the inverted reflectivity resolution and noise. Increasing λ decreases the resolution of inverted reflectivity, and decreasing λ may cause noise amplification. In practice, the proper λ value is data dependent and determined empirically.

Sparse-spike inversion (SSI) is an alternate minimal L_1 norm inversion method. Distinct from SSI, BPI utilizes the primal-dual log-barrier algorithm (Chen et al., 2001), whereas SSI solves only the primal equation (Taylor et al., 1979). Both BPI and SSI have a trade-off factor λ to control the inversion output. However, the values are not comparable between the methods. Both must be determined independently by test trials. This makes comparison of SSI and BPI results difficult, because the results are dependent on the parameterizations. In this study a wide range of λ s for each inversion type is used; the correlation of inverted reflectivity with true reflectivity is then determined for all λ s for both inversion methods, and the best BPI result is compared to the best SSI result. The

correlation coefficient between true and inverted reflectivities quantitatively combines the accuracy in the location of and relative magnitude of spikes.

To test and compare the methods, reflectivity series with randomly selected amplitudes and time-thicknesses were generated with a sample rate of 1 ms. Figures 2 and 3 show the SSI and BPI inversion results for a particular case. The reflection coefficient at each sample is a random value between -0.2 and 0.2 (shown in Figure 2b). Figure 2a shows the synthetic seismogram, which is the convolution of the reflectivity with a 40 Hz Ricker wavelet. Gaussian noise, giving a signal-to-noise ratio (S/N) of 10, which is generated by randomly drawing numbers from a Gaussian

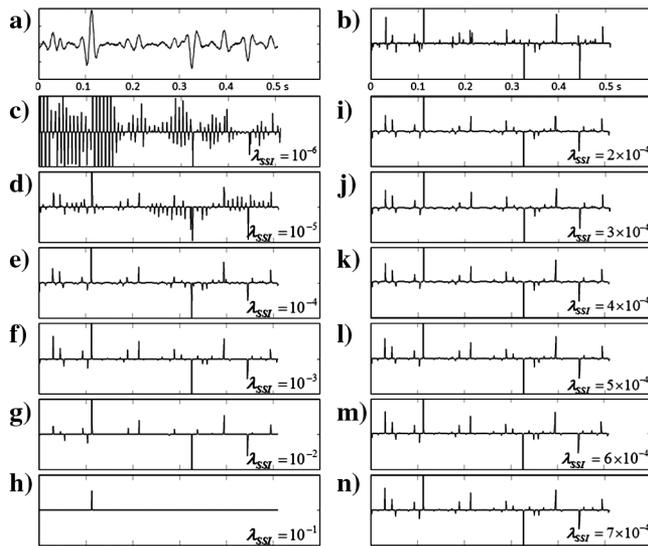


Figure 2. 1D synthetic tests of SSI are applied on (a) synthetic seismogram with 40 Hz Ricker wavelet and 10% random noise. (b) Shows the true reflectivity. (c)-(n) Shows SSI inversion results with varying λ_{SSI} values of (c) $\lambda_{SSI} = 10^{-6}$, (d) $\lambda_{SSI} = 10^{-5}$, (e) $\lambda_{SSI} = 10^{-4}$, (f) $\lambda_{SSI} = 10^{-3}$, (g) $\lambda_{SSI} = 10^{-2}$, (h) $\lambda_{SSI} = 10^{-1}$, (i) $\lambda_{SSI} = 2 \times 10^{-4}$, (j) $\lambda_{SSI} = 3 \times 10^{-4}$, (k) $\lambda_{SSI} = 4 \times 10^{-4}$, (l) $\lambda_{SSI} = 5 \times 10^{-4}$, (m) $\lambda_{SSI} = 6 \times 10^{-4}$; (n) $\lambda_{SSI} = 7 \times 10^{-4}$.

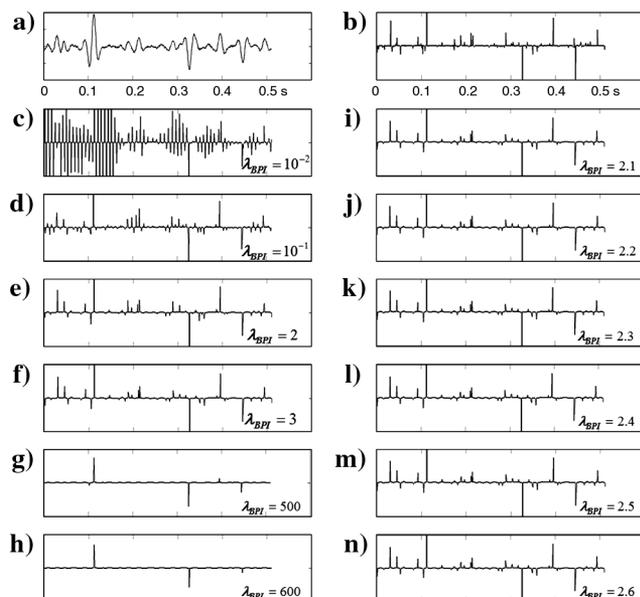


Figure 3. 1D synthetic tests of BPI are applied on (a) Synthetic seismogram with 40 Hz Ricker wavelet and 10% random noise. (b) True reflectivity. (c)-(n) Shows BPI inversion results with varying λ_{BPI} values of (c) $\lambda_{BPI} = 10^{-2}$, (d) $\lambda_{BPI} = 10^{-1}$, (e) $\lambda_{BPI} = 2$, (f) $\lambda_{BPI} = 3$, (g) $\lambda_{BPI} = 500$, (h) $\lambda_{BPI} = 600$, (i) $\lambda_{BPI} = 2.1$, (j) $\lambda_{BPI} = 2.2$, (k) $\lambda_{BPI} = 2.3$, (l) $\lambda_{BPI} = 2.4$, (m) $\lambda_{BPI} = 2.5$, and (n) $\lambda_{BPI} = 2.6$.

Figure 4. (a) λ -correlation curve for SSI based on the synthetic data in Figure 2. The correlation coefficient is calculated between SSI results and true reflectivity and is plotted against the regularization factor λ_{SSI} ; (b) the λ -correlation curve of BPI based on the same synthetic data in Figure 2. The correlation that is calculated between BPI results and true reflectivity is plotted with the regularization factor λ_{BPI} ; (c) and (d) show coarse λ -correlation curves of BPI and SSI based on the same synthetic data in Figure 2 for (c) noise-free and (d) 40% noise.

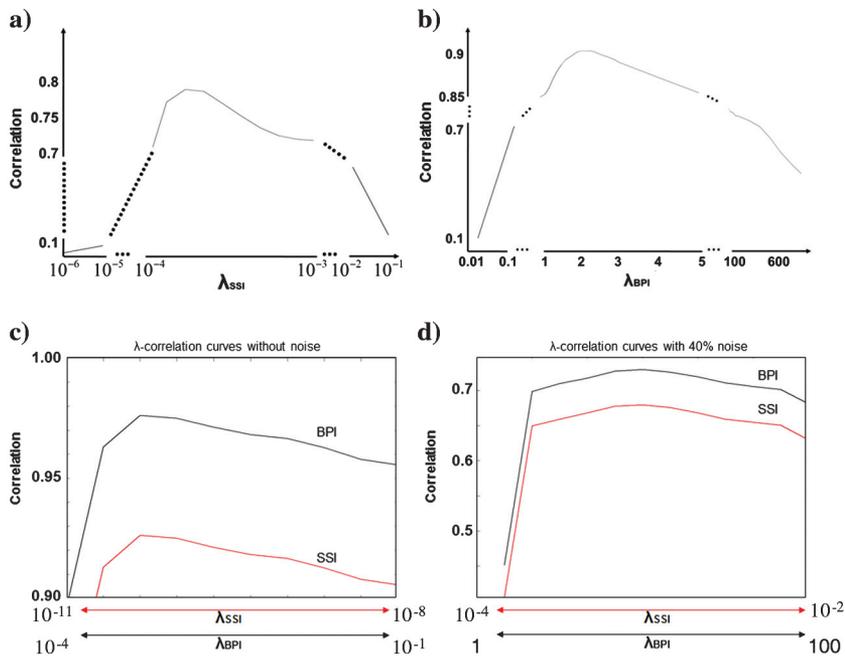


Figure 5. A crossplot of maximum correlation from BPI and SSI results calculated on 100 different randomly generated synthetic reflectivities. The dashed line represents equal correlation for both methods. BPI usually correlates better with the true reflectivity.

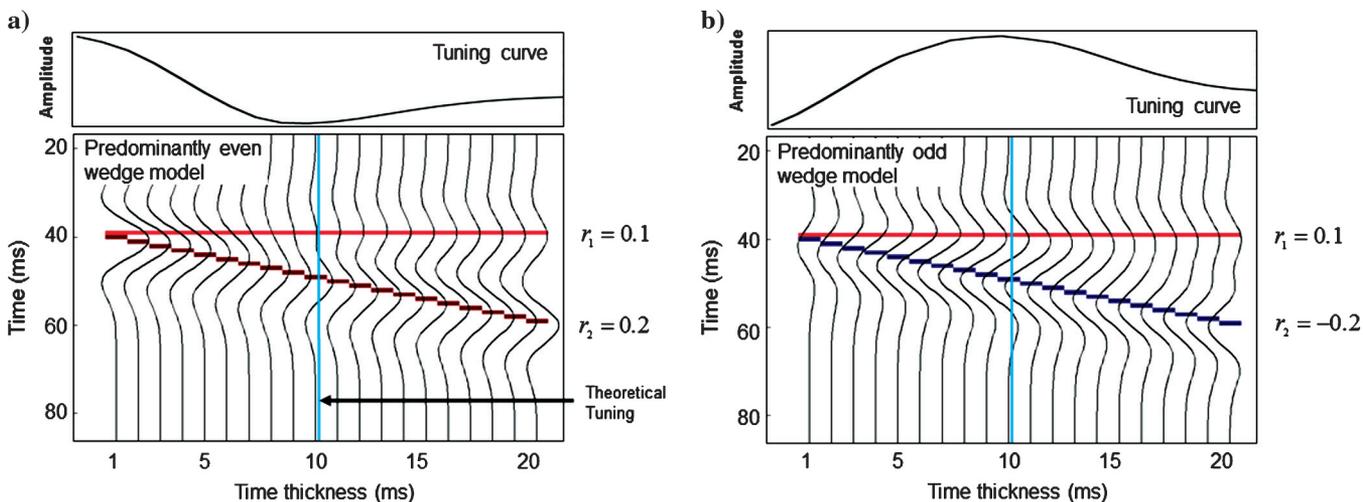
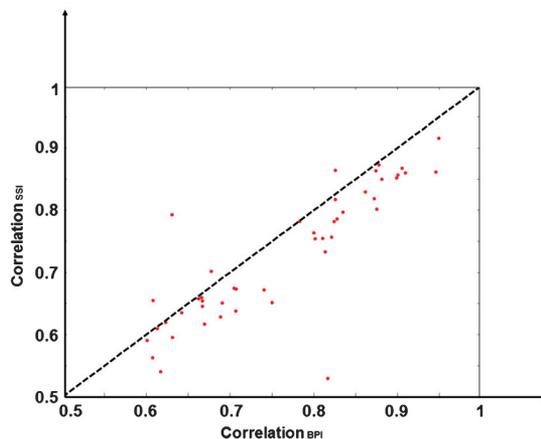


Figure 6. Predominantly even and odd wedge models. (a) A predominantly even wedge-model and its tuning curve with seismic responses overlaid. (b) A predominantly odd wedge-model and its tuning curve with seismic responses overlaid. Seismic responses are generated with a 40-Hz Ricker wavelet convolved with the wedge reflectivity. Tuning thickness is 10 ms.

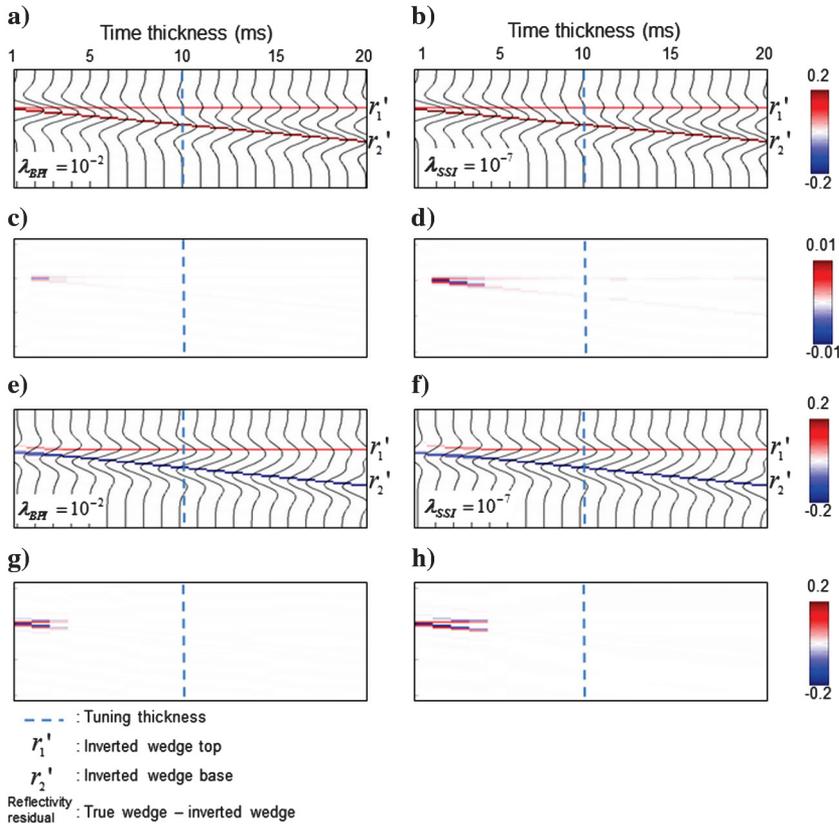


Figure 7. Predominantly wedge inversion without noise. (a)-(d) Shows tests results for a predominantly even wedge-model. (a) BPI inverted results with $\lambda_{BPI} = 10^{-2}$. (b) SSI inverted results with $\lambda_{SSI} = 10^{-7}$. (c) Residuals from BPI inverted results. (d) Residuals from SSI inverted results. (e)-(h) Shows tests results for predominantly odd wedge-model. (e) BPI inverted results with $\lambda_{BPI} = 10^{-2}$. (f) SSI inverted results with $\lambda_{SSI} = 10^{-7}$. (g) Residuals from BPI inverted results. (h) Residuals from SSI inverted results.

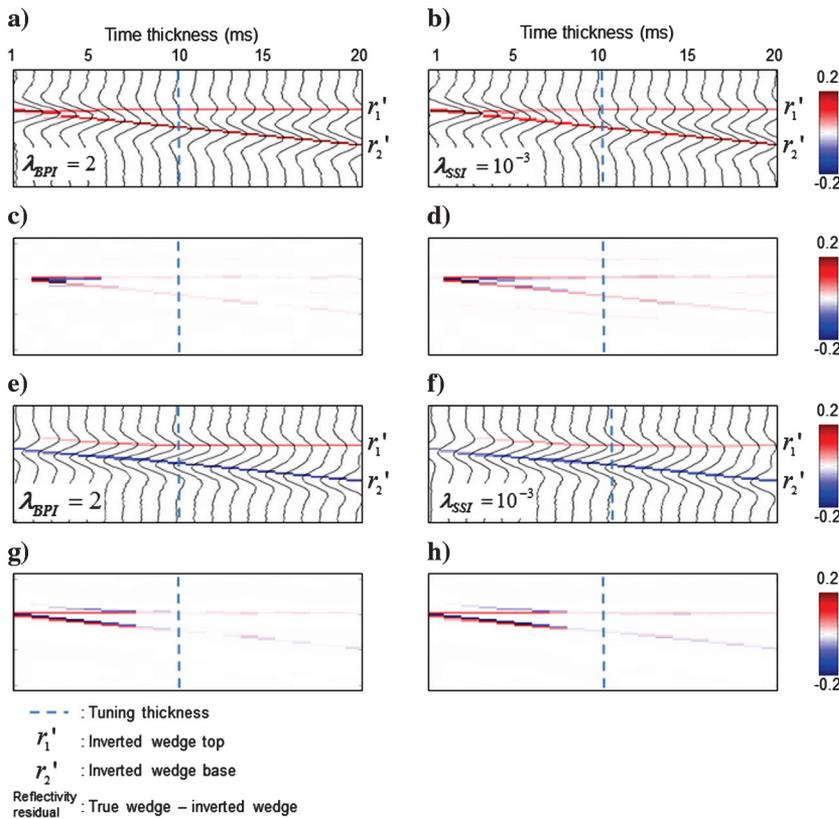


Figure 8. Predominantly wedge inversion with 10% noise. Letters (a) through (d) show tests results for a predominantly even wedge-model with 10% noise. (a) BPI inverted results with $\lambda_{BPI} = 2$. (b) SSI inverted results with $\lambda_{SSI} = 10^{-3}$. (c) Residual from BPI inverted results. (d) Residual from SSI inverted results. Letters (e) through (h) show tests results for a predominantly odd wedge-model with 10% noise. (e) BPI inverted results with $\lambda_{BPI} = 2$. (f) SSI inverted results with $\lambda_{SSI} = 10^{-3}$. (g) Residual from BPI inverted results. (h) Residual from SSI inverted results.

distribution with the L_2 norm that is 10% that of the L_2 norm of the noise-free synthetic seismogram, is added to the noise-free synthetic seismogram. Small λ_{SSI} values boost the noise on the inverted reflectivities (Figure 2c and 2d) and large λ_{SSI} values produce inverted reflectivities that are too sparse (Figure 2g and 2h). Extreme values of λ_{SSI} result in a discrepancy between the inverted and true reflectivity. For this synthetic example, the best inversion results are for regularization parameters between 10^{-4} and 10^{-3} . Further tests are performed with λ_{SSI} between 10^{-4} and 10^{-3} (Figure 2i–2n).

Figure 3 shows the BPI inverted reflectivities based on the same synthetic data in Figure 2a and 2b. Similar to SSI, a small λ_{BPI} causes noise to amplify (Figure 3c and 3d), and a large λ_{BPI} produces sparse results (Figure 3g and Figure 3h). The maximum correlation is between two and three (Figure 3i–3n).

Zero-lag correlations between the inverted and true reflectivities from previous tests have been plotted in Figure 4. Figure 4a shows that the maximum correlation for SSI tests is around 0.8 with λ_{SSI}

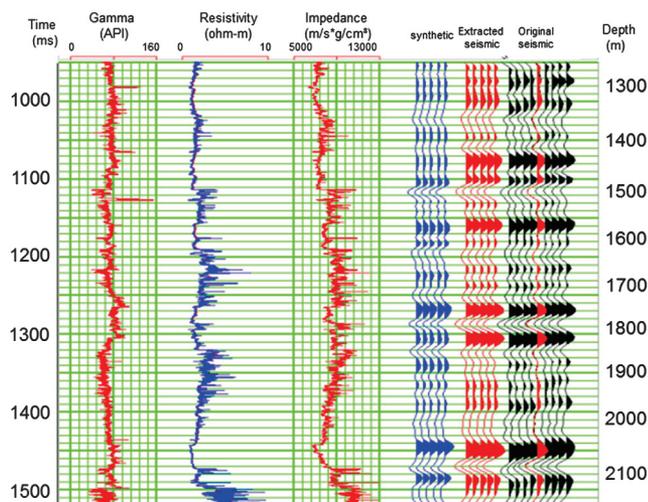


Figure 9. Seismic well tie, including gamma ray, resistivity, and computed impedance, along with the synthetic seismogram (blue), composite seismic trace at the well location (red), and original seismic traces near the well. The correlation coefficient between the synthetic and the composite trace is 0.79. Shale is indicated by higher gamma ray values.

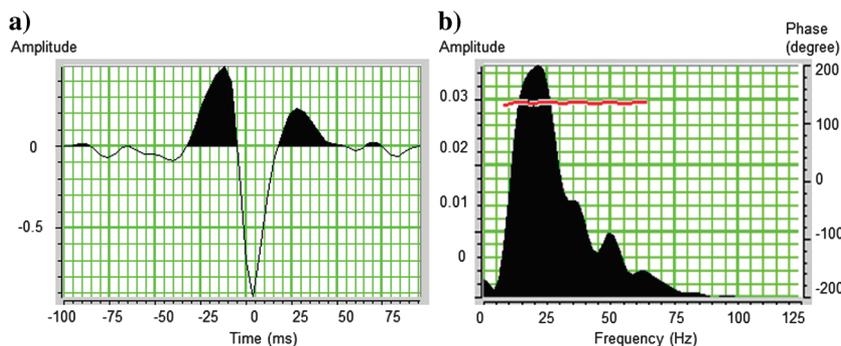


Figure 10. (a) Extracted seismic wavelet at well location using amplitude spectrum of seismic data averaged over 1000 ms and a spatial window consisting of 100 inlines and crosslines in the vicinity of the well, and application of constant phase rotations that produce the best synthetic tie. (b) Amplitude and phase spectrum of the wavelet.

around 2×10^{-4} . Figure 4b shows that the maximum correlation for BPI tests reaches about 0.9 at λ_{BPI} around 2.2.

In addition to the 10% noise level, synthetic tests were conducted for BPI and SSI on the same reflectivity series but with 0% and 40% noise levels (see Figure 4c and 4d). As expected, noise reduces the correlation coefficient for both cases, but BPI is degraded by a smaller percentage than SSI for this case.

The same methodology was applied using 100 different random reflectivity series and corresponding synthetic seismograms with a 40 Hz Ricker wavelet. For these synthetic cases, the noise level varied from 10% to 50% in increments of 10% with 20 traces at each noise level. In most cases, BPI correlates better to the true reflectivity model (Figure 5).

Resolution of the methods can be compared using wedge models. Wedge models may be primarily even or primarily odd (Figure 6). The tuning thickness t_R of a thin-bed model with a Ricker wavelet is given by Chung and Lawton (1995):

$$t_R = \frac{\sqrt{6}}{2\pi f_0}, \quad (11)$$

where f_0 is the dominant wavelet frequency. For a 40-Hz Ricker wavelet, $t_R = 10$ ms.

Figure 7a to 7d shows the best noise-free inversion results for the predominately even wedge-model for both BPI and SSI. For noise-free data, BPI has smaller residuals, especially for thin beds. Notice that the residuals are small compared to the data. Figure 7e–7h shows the best BPI and SSI results for the predominately odd wedge-model. For thin beds, the residual is of the same order of magnitude as the data. Overall, for noise-free data, BPI has a better resolution and smaller residual than SSI.

Figure 8a–8d shows the BPI and SSI inversion results of a predominately even wedge-model with a S/N of 10. The resolution of the two methods is comparable, with BPI being slightly better. The residuals are of the same order of magnitude as the data. BPI has better residuals than SSI. Figure 8e–8h shows the BPI and SSI inversion results of the predominately odd wedge-model with an S/N of 10. Resolution and residuals are slightly better for BPI than for SSI.

This series of synthetic tests suggests that sparse-layer BPI, using a wedge-model dictionary, is at least as good as SSI to invert for blocky earth models containing thin layers and usually better in

regards to resolution and fit to the data, when optimal regularization parameters are chosen. It remains to be seen, in practice, if this benefit can be realized, because the optimal regularization parameter may not be known. Nevertheless, the results are encouraging and suggest that even greater improvement could be obtained by further constraining the inversion using actual well-log reflectivity patterns; this is left as an objective of future research.

REAL DATA RESULTS

The BPI inversion with a wedge-model is tested on a 3D seismic data set in a clastic basin. Figure 9 shows the synthetic tie between the input seismic data and the well-log data. The extracted wavelet used for the synthetic

(Figure 10) is calculated from an inversion in which the wavelet is the unknown to be determined from the known well-log reflectivity and seismic traces at the well-log location. The peak frequency of the extracted wavelet in this case is close to 25 Hz (Widess, 1973, 1985), yielding a one-half-period time-thickness resolution of about 20 ms. The phase of the wavelet is about 140°.

The seismic well tie (Figure 9) shows a good fit ($r = 0.79$) between the unedited reflectivity convolved with the extracted wavelet and the seismic trace. However, because seismic data are much lower in frequency than log data, the fit is useful only as an approximation for aligning gross lithologic packages.

In the examples presented, both BPI and SSI were run using λ values that resulted in the best correlation to the well-log. BPI reflectivity inverted from the original seismic data is converted to band-limited impedance by the standard method of integration followed by low-cut filtering to avoid the systematic accumulation of errors (Figure 11).

As is evident in Figure 11, in the original seismic data it is almost impossible to detect or resolve thin layers, because some are apparent on the inverted impedance section. One of these thin layers is highlighted in Figure 11, in which a shale, embedded in thick sand, is indicated on the gamma ray log. This shale could be an important vertical flow barrier. Although it is not mappable on the original seismic data, it appears to have been detected after BPI inversion for band-limited impedance. The shale thickness, which is about 3 ms, is well below the tuning thickness (the peak frequency of the data is 25 Hz, yielding a one-quarter wavelength resolution of about 20 ms). The band-limited impedance from BPI appears to detect and resolve some, but not all, thin layers seen in the logs. Such detection could be fortuitous; however, the odds of detecting more than one thin anomalous layer on the inverted seismic at the same record times that the layers occur in the well logs is low enough that such a correlation cannot be discounted. In application,

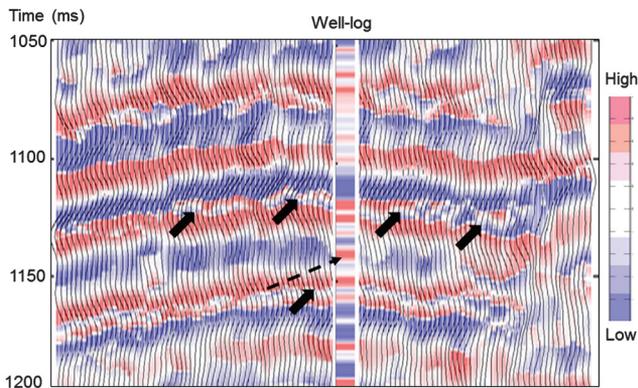


Figure 11. Inverted band-limited seismic impedance section (colors), inserted well-log band-limited impedance (central color bar) and original seismic wiggles. Red indicates high impedance and blue indicates low impedance. Solid arrows indicate the expression of thin low impedance shales seen in the well logs on the inverted impedance section. Not all thin layers evident on the well logs are resolved by the inversion. The dashed arrow indicates a relatively high impedance layer seen in the well-log but not in the seismic inversion at the well location, suggesting multiple solutions with similar seismic responses. Thin layers evident in the well-log are sometimes but not always apparent on the seismic inversion, suggesting data and case dependent resolution. None of the layers is resolved on the original seismic data.

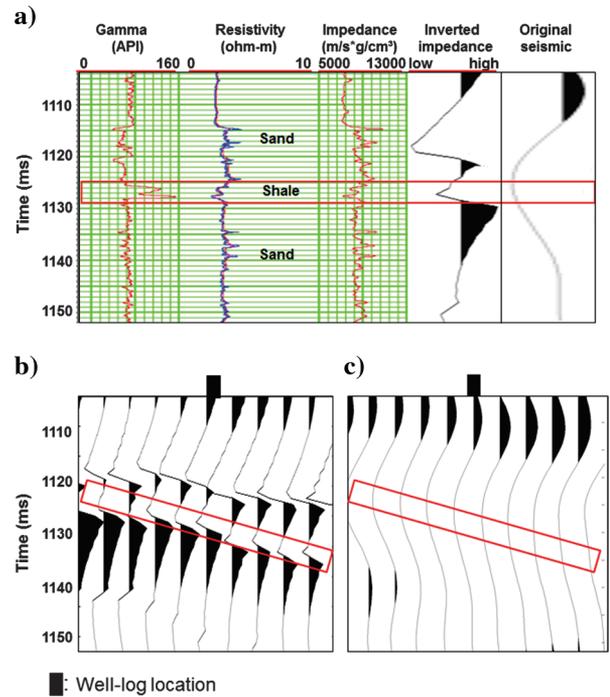


Figure 12. Zoom of well logs and seismic data in the vicinity of a thin shale highlighted by four solid arrows in Figure 11 at about 1130 ms. (a) Well logs. (b) Inverted band-limited seismic impedance using BPI and $\lambda_{BPI} = 2$. (c) Original seismic data. The red boxes show the interpreted location of the shale on the well logs and on the seismic data. The location of the well on the seismic data is indicated by the black box.

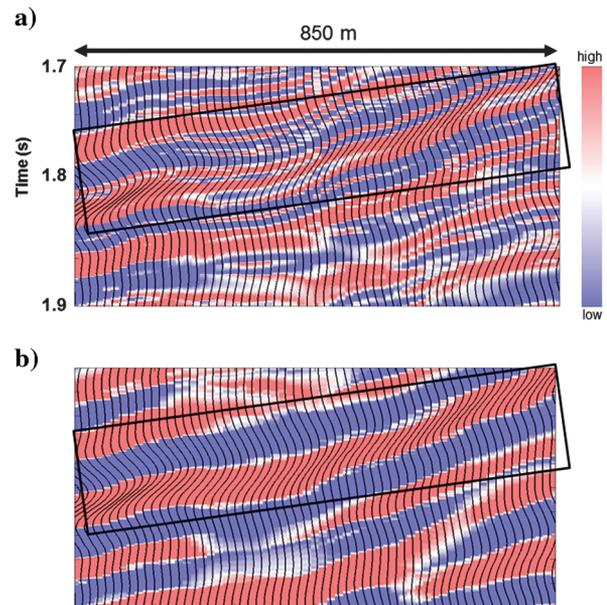


Figure 13. (a) BPI inverted band-limited impedance suggests many thin layers within the rectangle. (b) SSI inverted data shows a low to high impedance transition within the rectangle. The original seismic image data overlain as wiggles shows one layer of positive continuous amplitudes that can be picked as horizon. Red indicates high band-limited impedance and blue indicates low band-limited impedance.

more wells would be needed for validation before one could be confident that a particular layer was being detected and resolved. This example suggests that sparse-layer inversion provides a different and possibly significantly better representation of the layering observed in the log data than the original seismic image.

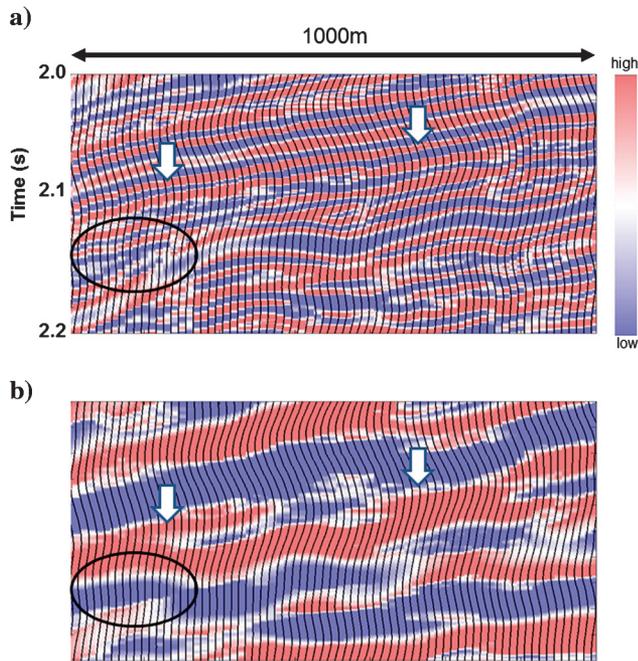


Figure 14. (a) BPI result reveals couples of laterally continuous dipping layering (white arrow) characteristic of undisturbed layer-cake geology. (b) SSI result shows a much simpler impedance layering structure. The black ellipse in (a) emphasizes a wedge structure inverted by BPI that is unclear in the seismic wiggles or the SSI.

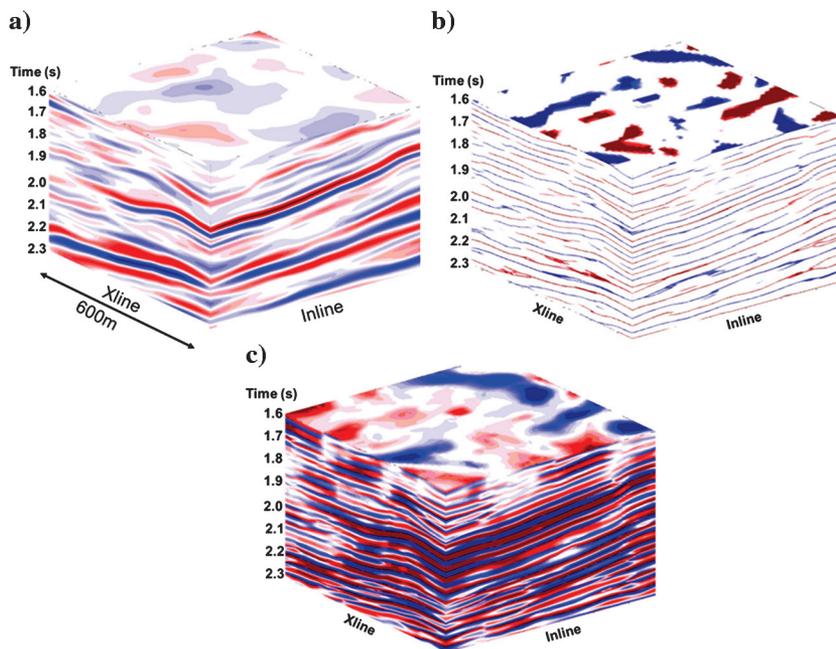


Figure 15. (a) Original 3D data set cube; (b) BPI inverted reflection coefficients cube; (c) BPI inverted band-limited impedance.

In practice, such an apparent increase in resolution could be due to the simple magnification of noise resulting from the selection of a regularization parameter that is too small, and comparison of the results to existing well logs and stratigraphic interpretations is needed to validate the resolution of the output. A regularization parameter λ_{BPI} of two was used to generate the result shown in Figures 11 and 12. If sufficient well control is available, one can optimize the regularization parameter to give inversion results that best match all well information. Unfortunately, no other wells were available in this case. In this example, as would often be the case in practice, parameterization is a judgment call that can greatly impact the validity of the results.

Figures 13 and 14 compare BPI and SSI inverted band-limited impedances with the original seismic wiggles overlain for many representative sections from a deeper interval in the 3D data set. The calibrated wavelet has a peak frequency of about 10 Hz and a 150° phase with the tuning time-thickness for this data being 50 ms. BPI clearly produces laterally consistent thin layers that are not revealed on the original seismic data or by SSI (Figure 13).

In Figure 14, although the BPI inverted band-limited impedance image is much higher in frequency than the seismic data or the SSI result, it exhibits more lateral coherence than the low frequency images, while also revealing a thinning wedge that is not apparent on the other images. The lateral coherence of the output of the BPI inversion process, which is a trace independent process with no lateral continuity constraints, spatial filtering, or starting model to impose false lateral continuity, is an important attribute that cannot be ignored. This observation was also made by [Puryear and Castagna \(2008\)](#) for spectral inversion results. They hypothesized that such a lateral consistency of the inversion result is caused by the reduction of wavelet side-lobe interference effects by the inversion process. A 3D minicube from the same data set (Figure 15) shows good lateral continuity of BPI inverted reflection coefficients and band-limited impedance.

DISCUSSIONS

The increased frequency content of the output of the sparse-layer inversion relative to the original seismic data is a consequence of the assumptions built into the process: (1) each reflection event is produced by the superposition of a limited number of layer responses, (2) the impedance structure of the earth is blocky, rather than transitional, and (3) there is a predefined range of layer thicknesses. To the extent that these assumptions are correct, a significant valid bandwidth extension can be achieved. If these assumptions are not a good representation of the earth structure, the inversion output will be incorrect.

In this study, a constant wavelet is utilized for simplicity. In practice, a set of time and spatial varying wavelets could improve the method. Such a set of wavelets can be easily implemented by modifying the wavelet kernel matrix, but it is difficult to get a reliable time and spatial varying trends. In this case, attenuation, dispersion, and many other wave propagation effects have to be considered.

In our formulation for BPI inversion, we assume horizontal isotropic homogenous layers, a known wavelet, and that the convolutional model is applicable. These assumptions are clearly wrong; yet, such assumptions have underlain many processing and inversion methods that have been used successfully in commercial applications. Part of the utility of these assumptions is related to the fact that seismic data is interpreted largely in a relative sense, and when absolute numbers are produced much calibration and correction is applied to the output. As we push resolution beyond that of the original seismic data we expect such assumptions to become even more important. It remains to be seen how successful BPI inversion will be in predicting absolute rock properties. It is probably safer at this point to view the output of BPI inversion simply as another seismic attribute that can be used for visualization purposes and for multiattribute analysis.

For data with a high S/N and significant asymmetry in the local impedance structure (leading to an appreciable even component of the reflectivity), thicknesses far below tuning can potentially be resolved given sufficient S/N. Noise in the data deteriorates the performance of the inversion. Predominantly odd reflectivity is less readily resolved than reflectivity with a large even component.

Sparse-layer inversion using basis pursuit of a wedge dictionary does not use well-log information in the inversion as a starting model or as a constraint. Therefore, the resulting inversion is unbiased by preconceived ideas, other than those of the general assumptions stated above.

Basis pursuit inversion can be readily extended to the prestack AVO domain. One approach would be to invert for reflection coefficients on angle stacks. Alternatively, one could formulate the inversion using linearized Zoeppritz approximations (e.g., Aki and Richards, 1980), and invert the angle gather simultaneously for a spiky series of $\Delta V_p/V_p$, $\Delta V_s/V_s$, and $\Delta\rho/\rho$ (or similar) AVO attributes. One would have to take special care in accounting for the offset dependence of the seismic wavelet. Other difficulties to consider would be the lower S/N of the prestack data relative to the stack and the greater impact of moveout error at higher frequencies. The study of prestack BPI is left as an objective of future research.

Although we have not optimized the inversion code, the method is obviously computationally more intensive than sparse-spike inversion, because, rather than inverting for the coefficients of the seismic wavelet, an entire dictionary of dipole responses must be considered in lieu of a single wavelet. However, the problem is small enough that we have been able to implement all the work presented here on a desktop PC. Furthermore, being a trace-by-trace process, the method is readily parallelizable and is well suited to implementation on a cluster.

CONCLUSIONS

Sparse-layer inversion can be accomplished by basis pursuit of a dictionary of functions, representing thin-bed reflectivity patterns, and the constitution of the seismic trace as a superposition of these patterns. This method determines a sparse number of patterns summed together to form the seismic trace. Synthetic tests indicate that sparse-layer inversion, using basis pursuit (BPI), can better resolve thin beds than a comparable sparse-spike inversion (SSI) and usually correlates better to known reflectivity when optimal regularization parameters are used for both methods.

We conclude that the sparse-layer inversion of seismic reflection data using basis pursuit with a wedge-model dictionary is a viable

inversion method. The results are comparable to and, arguably, more accurate and more highly resolved than sparse-spike inversion. Most significantly, the inversion results produce different images that lead to different geological interpretations of reflection events than conventional data or sparse-spike inversion.

The application of sparse-layer inversion to field data suggests but does not prove improved detectability, resolution, and lateral continuity of thin layers. In one instance highlighted here, a 3 ms thick shale not visible on the original seismic data was detected and resolved. In practice, one needs to determine whether such apparent resolution is real or a consequence of using too small a regularization parameter; the best proof would be obtained by comparing the inversion results to multiple well logs, where there is a significant lateral variation in thickness of the layer. Commonly, as was the case here, only limited well information is available, and confidence must be built based on how well the image conforms to geological expectations. Comparison of inversion images to seismic and conventional sparse-spike sections shows alternative definition of what appear to be stratigraphic features. Lateral continuity of BPI inversion images can be greater than the continuity of the original seismic data, suggesting that wavelet side-lobe interference that produces apparent discontinuities is reduced by the inversion process. As evidenced by synthetic and real data results, BPI can potentially be a practical tool for seismic exploration and reservoir characterization purposes.

REFERENCES

- Aki, K., and P. G. Richards, 1980, *Quantitative seismology: Theory and methods*: W. H. Freeman, A Series of books in geology.
- Bork, J., and L. Wood, 2001, Seismic interpretation of sonic logs: 71st Annual International Meeting, SEG, Expanded Abstracts, 510–513.
- Chen, S. S., D. L. Donoho, and M. A. Saunders, 2001, Atomic decomposition by basis pursuit: *Society for Industrial and Applied Mathematics Review*, **43**, no. 1, 129–159.
- Chopra, S., J. Castagna, and O. Portniaguine, 2006, Thin-bed reflectivity inversion: 76th Annual International Meeting, SEG, Expanded Abstracts, 2057–2061.
- Chopra, S., J. Castagna, and Y. Xu, 2009, Thin-bed reflectivity inversion and some applications: *First Break*, **27**, 17–24.
- Chung, H., and D. C. Lawton, 1995, Frequency characteristics of seismic reflections from thin beds: *Canadian Journal of Exploration Geophysics*, **31**, 32–37.
- Cooke, D. A., and W. A. Schneider, 1983, Generalized linear inversion of reflection seismic data: *Geophysics*, **48**, 665–676.
- Debye, H. W. J., and P. van Riel, 1990, LP-norm deconvolution: *Geophysical Prospecting*, **38**, 381–404.
- Hu, C., P. Stoffa, and K. McIntosh, 2008, First arrival stochastic tomography: Automatic background velocity estimation using beam semblances and VFSA: *Geophysical Research Letters*, **35**, 3307–3311.
- Kaarensen, K., and T. Taxt, 1998, Multichannel blind deconvolution of seismic signals: *Geophysics*, **63**, 2093–2107.
- Kormylo, J., and J. M. Mendel, 1980, Maximum-likelihood seismic deconvolution: 50th Annual International Meeting, SEG, Session: G.77.
- Mosegaard, K., and P. D. Vestergaard, 1991, A simulated annealing approach to seismic model optimization with sparse prior information: *Geophysical Prospecting*, **39**, 599–612.
- Nguyen, T., and J. Castagna, 2010, High resolution seismic reflectivity inversion: *Journal of Seismic Exploration*, **19**, 303–320.
- Oldenburg, D. W., T. Scheuer, and S. Levy, 1982, Recovery of acoustic impedance from reflection seismograms: 52nd Annual International Meeting, SEG, Session: S15.4.
- Oldenburg, D. W., T. Scheuer, and S. Levy, 1983, Recovery of the acoustic impedance from reflection seismograms: *Geophysics*, **48**, 1318–1337.
- Partyka, G., 2005, Spectral Decomposition, Spring 2005 SEG Distinguished Lecturer.
- Porsani, M. J., and B. Ursin, 2000, Mixed-phase deconvolution and wavelet estimation: *The Leading Edge*, **19**, 76–79.
- Portniaguine, O., and J. Castagna, 2005, Spectral inversion: Lessons from modeling and boonesville case study: 75th Annual International Meeting, SEG, Expanded Abstracts, **24**, 1638–1641.

- Purveyar, C. I., and J. P. Castagna, 2008, Layer-thickness determination and stratigraphic interpretation using spectral inversion: *Geophysics*, **73**, no. 2, R37–R48.
- Riel, P. V., and A. J. Berkhout, 1985, Resolution in seismic trace inversion by parameter estimation: *Geophysics*, **50**, 1440–1455.
- Sacchi, M. D., D. R. Velis, and A. H. Cominguez, 1994, Minimum entropy deconvolution with frequency-domain constraints: *Geophysics*, **59**, 938–945.
- Tarantola, A., 2004, *Inverse problem theory and methods for model parameter estimation*: Society for Industrial and Applied Mathematics.
- Taylor, H. L., S. C. Banks, and J. F. McCoy, 1979, Deconvolution with the ℓ_1 norm: *Geophysics*, **44**, 39–52.
- Ulrych, T. J., D. R. Velis, and M. D. Sacchi, 1995, Wavelet estimation revisited: *The Leading Edge*, **14**, 1139–1143.
- Wang, Y., 2007, Seismic time-frequency spectral decomposition by matching pursuit: *Geophysics*, **72**, no. 1, V13–V20.
- Wang, Y., 2010, Multichannel matching pursuit for seismic trace decomposition: *Geophysics*, **75**, no. 4, V61–V66.
- Widess, M. B., 1973, How thin is a thin bed?: *Geophysics*, **38**, 1176–1180.
- Widess, M. B., 1985, How thin is a thin bed?: *Geophysics*, **50**, 2061–2065.
- Zhang, R., 2010, *Basis pursuit seismic inversion*: Ph.D. thesis, University of Houston.