GEOPHYSICS, VOL. 69, NO. 2 (MARCH-APRIL 2004); P. 608–616, 10 FIGS.
10.1190/1.1707081

Interpolation of bathymetry data from the Sea of Galilee: A noise attenuation problem

Antoine Guitton* and Jon Claerbout*

ABSTRACT

We process a bathymetry survey from the Sea of Galilee. This data set is contaminated with non-Gaussian noise in the form of spikes inside the lake and at the track ends. There is drift on the depth measurements leading to vessel tracks in the preliminary depth images. The drift comes from different seasonal and human conditions during data acquisition, e.g., wind and water levels. We derive an inversion scheme that produces a much-reduced noise map of the Sea of Galilee. This inversion scheme includes preconditioning and iteratively reweighted least squares with the proper weighting function to remove the non-Gaussian noise. We remove the ship tracks by adding a modeling operator inside the inversion that accounts for the drift in the data. We then approximate the model covariance matrix with a prediction error filter that enhances details inside the lake. Unfortunately, the prediction error filter slightly degrades the frequency content of the final depth map. Our images of the Sea of Galilee show ancient shorelines and, inside the lake, rifting features.

INTRODUCTION

In this problem, we are given depth sounding data from the Sea of Galilee. The Sea of Galilee is unique because it is a freshwater lake below sea level. It seems to be connected to the Great Rift (pull-apart) Valley crossing East Africa. The ultimate goal is to produce a good map of the depth to bottom and images useful for identifying archaeological, geological, and geophysical details of the sea bottom. In particular, we hope to identify some ancient shorelines around the lake and meaningful geological features inside the lake. The ancient shorelines could unravel early settlements of archeological interest or old fishing ports.

The raw data (Figure 1), irregularly distributed across the surface, are 132 044 triples (x_i, y_i, z_i), where x_i ranges over about 12 km and where y_i ranges over about 20 km. We want to interpolate the data to a regular grid using inversion to facilitate the processing, such as noise removal, and to create a map that can be easily analyzed for identifying artifacts and geology. The pertinence of this data set to our daily geophysical problems is threefold. First, we often have to interpolate seismic maps (Britze, 1998), potential field data (Guspi and Introcaso, 2000), or other measurements to compensate for the sparseness and irregularities of acquisition geometries. Second, as seen in the raw data in Figure 1, some noise bursts related to spurious electronic signals (glitches) and/or positioning errors need to be accounted for in the inversion scheme. This problem is common, for example, in tomography (Bube and Langan, 1997), deconvolution of noisy data (Chapman and Barrodale, 1983), and velocity analysis (Guitton and Symes, 2003), where outliers can degrade the final model if we assume a Gaussian distribution of the noise. Third, the final image of the Sea of Galilee will display the vessel tracks because the measurements on the lake were made on different days with different weather and human conditions. We can directly link this problem to the goal of removing the acquisition footprint with 3D seismic data (Duijndam et al., 2000; Schuster and Liu, 2001; Chemingui and Biondi, 2002). Therefore, the interpolation of the data from the Sea of Galilee becomes a problem of spiky noise and coherent noise attenuation. We solve the problem with a finely tuned inversion scheme that should work for other geophysical applications.

Many students at the Stanford Exploration Project have attempted to produce a satisfying map of the sea bottom. Fomel and Claerbout (1995) introduce the ℓ^1 norm via iteratively reweighted least squares (IRLS) to eliminate the noise bursts present in the data. Brown (2001) attempts to remove acquisition tracks by estimating the systematic error between tracks at crossing points. Karpushin and Brown (2001) use a bank of prediction error filters (PEFs) to whiten the residual along tracks. However, in most of these results, a loss of resolution hampers the goal of seeing small features in the final image. In this paper, we borrow ideas from these authors with three new approaches. The first important twist is preconditioning

Published on Geophysics Online October 8, 2003. Manuscript received by the Editor January 8, 2003; revised manuscript received October 2, 2003.
*Stanford University, Stanford Exploration Project, Geophysics Department, Mitchell Building, Stanford, California 94305-2215. E-mail: antoine@sep.stanford.edu, jon@sep.stanford.edu.

(Fomel, 2001), the second is modeling the ship's track instead of filtering it from the residual, and the third is estimating the model covariance with a PEF. We show that preconditioning with an IRLS method removes the glitches and noise bursts very well. In addition, modeling the ship tracks within the inversion almost entirely removes the acquisition footprint. The PEF unravels small details on the shores of the Sea of Galilee.

ATTENUATION OF NOISE BURSTS AND GLITCHES

Now we show our formulation of the regridding problem. Let $\mathbf{h}$ be an abstract vector containing as components the water depth over a 2D spatial mesh, and let $\mathbf{d}$ be an abstract vector whose successive components represent depth along the vessel tracks. One way to grid irregular data is to minimize the length of the residual vector $\mathbf{r}_d(\mathbf{h})$:

$$\mathbf{0} \approx \mathbf{r}_d = \mathbf{B}\mathbf{h} - \mathbf{d}, \tag{1}$$

where $\mathbf{B}$ is a 2D linear interpolation operator and $\mathbf{r}_d$ is the data residual. Note that sinc or spline interpolants could be used as well, but we opted for a simple linear interpolation operator to focus our analysis on the noise attenuation problem only. This fitting goal requires that the gridded data $\mathbf{h}$ take on appropriate values where the data $\mathbf{d}$ were collected. The bin size is 60×50 m so that the number of data points per bin is roughly constant and the aspect ratio of the lake is preserved in the number of samples in the vertical and horizontal directions. We display a simple binning (without interpolation or inversion) of the raw data (Figure 1) in Figure 2. Note that north points to the top of the map. A coarser mesh would avoid the empty bins but would lose resolution. As we refine the mesh for more detail, the number of empty bins grows as does the care needed in devising a technique for filling them. The black lines in Figure 2 are the ship tracks. Notice that some data points are outside the contour of the water. These must represent navigation errors. Figure 3 displays the ship tracks only. The straight lines in the north part of the lake result from positioning errors. The tracks match almost exactly the black lines in Figure 2.

Unless data are collected everywhere on a very fine mesh, and depending on how we parameterize the grid, the regridding may leave holes on the mesh. We can eliminate the holes by adding some regularization, such as

$$\begin{aligned} \mathbf{0} &\approx \mathbf{r}_d = \mathbf{B}\mathbf{h} - \mathbf{d}, \\ \mathbf{0} &\approx \epsilon \mathbf{r}_h = \epsilon \nabla \mathbf{h}, \end{aligned} \tag{2}$$

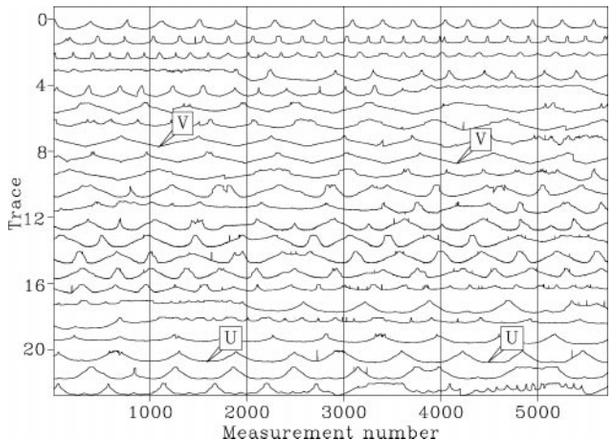

Figure 1. Depth of the Sea of Galilee along the vessel tracks. On one traverse across the lake, the depth record is *u* shaped. A few *v*-shaped tracks result from vessel turn-arounds. All points used for building the maps are shown here. Each point corresponds to one depth measurement inside the lake. The input data are displayed as seismic traces for display purpose only. The long signal is broken into 23 strips of 5718 depth measurements. The horizontal axis corresponds to a measurement number. The depth appears as one long track, although the surveying was done in several episodes that do not always continue the same track.

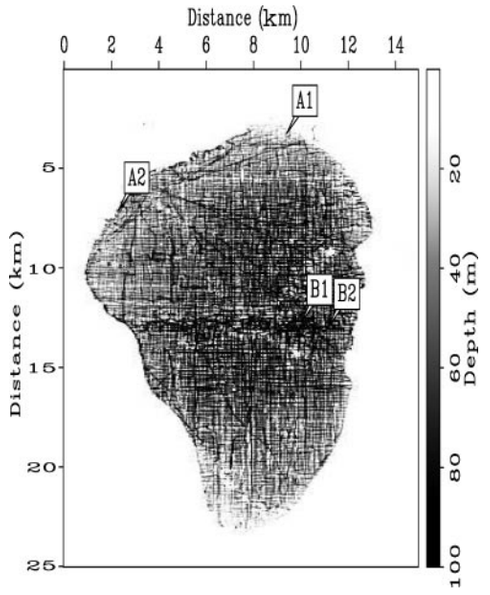

Figure 2. Simple binning of the raw data in Figure 1. The ship tracks and empty bins are visible and need to be accounted for in the inversion process. A1, B1, A2, and B2 point to track locations used in Figures 8 and 9. North at the top.

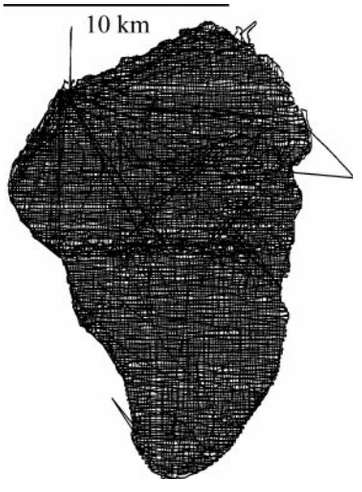

Figure 3. Ship tracks for the Sea of Galilee data set. The north part of the lake (top) has many navigation glitches which show up as long, straight lines. No track goes all the way from the north (top) to the south (bottom). Most of the tracks stop in the middle of the lake.

where $\nabla = (\partial/\partial x, \partial/\partial y)$ and $\mathbf{r}_h$ is the model space residual. We then minimize the misfit function,

$$f(\mathbf{h}) = \|\mathbf{r}_d\|^2 + \epsilon^2 \|\mathbf{r}_h\|^2, \tag{3}$$

to estimate the interpolated map of the lake. In theory (Tarantola, 1987), the regularization operator (squared) should be the model covariance operator, given an a priori model $\mathbf{h}_0$. Since we do not have an a priori model, we chose the gradient operator ∇ as a way of saying that the bottom of the lake is smooth. However, as pointed out by Harlan (1995), the regularization and the data fitting goal in equation (2) contradict each other. One equation tends to add details in the final map, whereas the second one (the regularization) tends to smooth it. We can more easily balance these two goals by preconditioning the problem (Fomel, 2001).

Preconditioning for accelerated convergence

A generally available preconditioning method is to change variables so that the regularization operator becomes an identity matrix (Claerbout and Fomel, 2002). The gradient ∇ in equation (2) has no inverse, but its spectrum $-\nabla'\nabla$, which appears in equation (3), can be factored ($-\nabla'\nabla = \mathbf{H}'\mathbf{H}$) into triangular parts $\mathbf{H}$ and $\mathbf{H}'$, where $\mathbf{H}$ is known as the helix derivative and ($'$) denotes the adjoint. This $\mathbf{H}$ is invertible by deconvolution (Claerbout, 1998). The fitting goals in equation (2) can be then rewritten

$$\begin{aligned} \mathbf{0} &\approx \mathbf{r}_d = \mathbf{B}\mathbf{H}^{-1}\mathbf{p} - \mathbf{d}, \\ \mathbf{0} &\approx \epsilon \mathbf{r}_p = \epsilon \mathbf{p}, \end{aligned} \tag{4}$$

with $\mathbf{p} = \mathbf{H}\mathbf{h} \approx \nabla \mathbf{h}$ and $\mathbf{r}_p$ the residual for the new variable $\mathbf{p}$. We then minimize the misfit function,

$$g(\mathbf{p}) = \|\mathbf{r}_d\|^2 + \epsilon^2 \|\mathbf{r}_p\|^2, \tag{5}$$

and finally compute $\mathbf{h} = \mathbf{H}^{-1}\mathbf{p}$ to estimate the interpolated map of the lake.

Experience shows that an iterative solution for $\mathbf{p}$ converges much more rapidly than an iterative solution for $\mathbf{h}$, thus showing that $\mathbf{H}$ is a good choice for preconditioning. There is no simple way of knowing beforehand the best value of ϵ. Practitioners like to see solutions for various values of ϵ which can be computationally expensive. Practical exploratory data analysis is more pragmatic: without a simple, clear theoretical basis, analysts generally begin from $\mathbf{p} = \mathbf{0}$ and then abandon the fitting goal $\mathbf{0} \approx \epsilon \mathbf{r}_p = \epsilon \mathbf{p}$ (Crawley, 2000; Rickett et al., 2001; Trad et al., 2003). Implicitly, they take $\epsilon = 0$. Then they examine the solution as a function of iteration, imagining that the solution at larger iterations corresponds to smaller ϵ and that the solution at smaller iterations corresponds to larger ϵ. In all our computations, we follow this approach and omit the regularization in the estimation of the depth maps.

ℓ^1 norm

We show how spikes and noise glitches can be attenuated with an approximate ℓ^1 norm. One main problem with the Galilee data is the presence of outliers in the middle of the lake and at the track ends. We could attenuate these spikes by editing or applying running median filters. However, the former involves human interpretation and the latter might compromise small details by flattening and distorting the signal (Claerbout and Fomel, 2002). Therefore, inversion appears to be the best compromise by eliminating the spikes while honoring the data in an automated way. Starting from equation (4), we can introduce a weighting operator that deemphasizes high residuals as follows:

$$\begin{aligned} \mathbf{0} &\approx \mathbf{r}_d = \mathbf{W}(\mathbf{B}\mathbf{H}^{-1}\mathbf{p} - \mathbf{d}), \\ \mathbf{0} &\approx \epsilon \mathbf{r}_p = \epsilon \mathbf{p}, \end{aligned} \tag{6}$$

with a diagonal matrix $\mathbf{W}$:

$$\mathbf{W} = \mathbf{diag}\left(\frac{1}{\left(1 + r_i^2/\bar{r}^2\right)^{1/4}}\right), \tag{7}$$

where r_i is the residual for one component of $\mathbf{r}_d$ and $\bar{r}$ is a prechosen constant. This weighting operator ranges from ℓ^2 to ℓ^1, depending on $\bar{r}$. It is somewhat difficult to evaluate a good $\bar{r}$ for a particular problem because the transition between the two norms is smooth. For this evaluation, we first selected a very small $\bar{r}$, tested it by running the inversion, and increased it by an order of magnitude. We repeated this loop four times, starting from $\bar{r} = 0.001$ cm. After this process, we selected $\bar{r} = 0.01$ cm, which is very small. This choice was not based on geophysical considerations; $\bar{r} = 10$ cm might appear a better choice since the measurements are recorded to an accuracy of about 10 cm. It is not based on statistical choices either, as done by Bube and Langan (1997). This choice of $\bar{r}$ simply gives us the most pleasing results after inversion, i.e., attenuation of all noise bursts in the final image while preserving the resolution and geological information of the final map. Because $\bar{r}$ is very small, we are essentially simulating an ℓ^1 norm only. The weighting operator $\mathbf{W}$ is kept constant for a number of iterations and is then reevaluated. The IRLS method is guaranteed to converge to the ℓ^1 estimate of the model parameters (Bube and Langan, 1997). The linear steps are computed with a conjugate gradient solver. Abandoning the damping in equation (6), i.e., $\epsilon = 0$, makes the IRLS method very appealing because we focus on minimizing the data residual. This is only possible by preconditioning the problem.

Now we test our proposed method to eliminate the outliers. We use a conjugate gradient solver to estimate the preconditioned variable $\mathbf{p}$, and we use the fitting goals in equations (4) and (6) to produce depth images of the Sea of Galilee. Equation (4) is referred as the ℓ^2 norm solution and equation (6) as the ℓ^1 norm solution.

In Figure 4a, we show $\mathbf{p}$ estimated with the ℓ^2 norm. We do not display the scale bar when we plot $\mathbf{p}$ because its values are of little interest for our analysis. Although $\mathbf{p}$ appears to be a variable of mathematical interest only, the solution $\mathbf{h}$ is so smooth that we have difficulty viewing it. We could view the two components of $\nabla \mathbf{h}$, but $\mathbf{p}$ is a roughened version of $\mathbf{h}$. In addition, it is more convenient to view $\mathbf{p}$ than the two images $\partial \mathbf{h}/\partial x$ and $\partial \mathbf{h}/\partial y$ because $\mathbf{p}$ is a single component vector. We can see considerable spurious noise everywhere on the map of Figure 4a. In addition, we can see the vessel tracks in the northern part. This first result is obtained after 1000 iterations, which means that we essentially simulate a least-squares solution without damping. Therefore, all data noise is inverted for and mapped

in the final model. With noisy data, common practice is to use a damped least-squares problem to minimize the effects of the noise. We can easily simulate a damped least-squares solution by decreasing the number of iterations, as explained earlier.

In Figure 4b, we show the roughened map of the Sea of Galilee after 50 iterations, thus recreating the solution of a damped least-squares problem. Most of the noise has been attenuated, but spikes are still present in the northern part of the lake (shown as O) so the ℓ^1 norm should be used.

Figure 4c displays $\mathbf{p}$ estimated with the ℓ^1 norm [i.e., equation (6) with a small $\bar{r}$]. Most of the glitches are attenuated, showing vessel tracks only. Some ancient shorelines in the western and southern parts of the Sea of Galilee are now easy to identify (shown as AS). In addition, we also start to see a valley in the middle of the lake (shown as R). This feature is also present in Figure 4b, where no attempts were made to remove the spikes. Therefore, this can be either a geological feature that represents the ongoing rifting in this area or a track. The next section will prove that this valley is not a processing artifact or some unaccounted noise in our inversion scheme. The data outside the sea have been also partially removed. The tracks (shown as T) are still clearly visible after the attenuation of the outliers because they do not fit the model of the noise we are trying to remove.

Figure 5 shows the bottom of the Sea of Galilee ($\mathbf{h} = \mathbf{H}^{-1}\mathbf{p}$) after inversion. Each line represents one east–west line of the interpolated data every 500 m. The ℓ^1 result is a great improvement over the ℓ^2 maps with or without damping. The glitches inside and outside the sea have disappeared. Also, the ℓ^1 norm gives positive depths everywhere. Although not visible everywhere in Figure 5, topography is produced outside the lake. Indeed, the effect of regularization is to produce synthetic topography, a natural continuation of the lake floor surface.

We have shown that a combination of preconditioning and IRLS removes the spikes in the depth map of the Sea of Galilee. In the next section, we propose removing the ship tracks by introducing an operator in equation (6) that models the coherent noise created by different weather and human conditions during data acquisition.

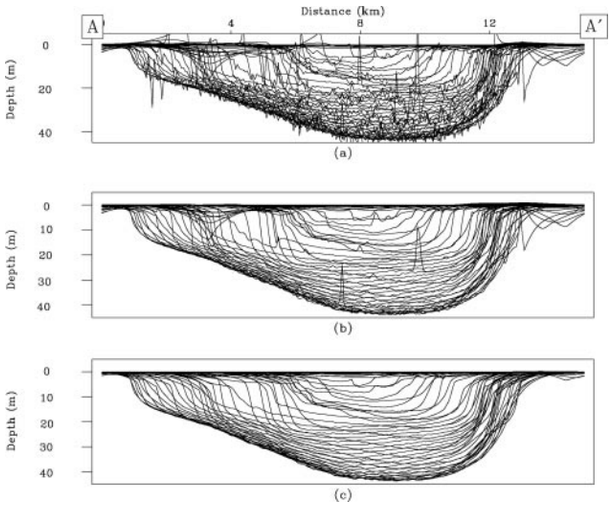

Figure 5. View of the bottom of the lake ($\mathbf{h} = \mathbf{H}^{-1}\mathbf{p}$) with the ℓ^2 norm after (a) 1000 iterations, which simulates a least-squares solution without damping, and (b) after 50 iterations, which simulates a least-squares solution with damping. (c) View of the bottom of the lake with the ℓ^1 norm. With the ℓ^1 norm, all depth values are now positive and the spikes have been attenuated. Each line represents one east–west track between A and A' in Figure 4a of the interpolated data every 500 m.

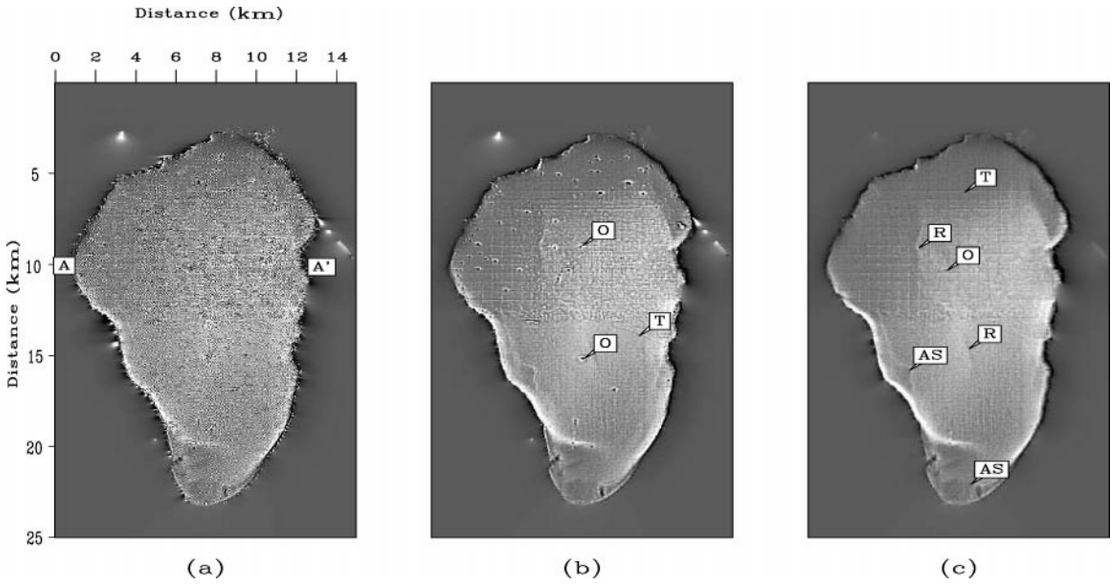

Figure 4. Estimated $\mathbf{p}$ with equation (4) in a least-squares sense after (a) 1000 and (b) 50 iterations, which simulates a least-squares solution (a) without and (b) with damping; $\mathbf{p}$ is a vector with only one dimension. The letters A and A' indicate the direction of the cross-sections in Figure 5. (c) Estimated $\mathbf{p}$ with equation (6) in an ℓ^1 sense. The spikes have been correctly attenuated. Some interesting features are shown by the arrows: AS points to few ancient shores, O points to some outliers, T points to few tracks, and R points to a ridge.

ATTENUATION OF SHIP TRACKS

We are now halfway to a noise-free image. As seen in Figure 4, the vessel tracks overwhelm possible small details at the bottom of the Sea of Galilee. By small details, we mean objects of 100 m and more, according to our bin size. In this section, we propose a strategy based on the assumption that the inconsistency between tracks comes mainly from different human and seasonal conditions during the acquisition. Unfortunately, we have no records of the weather and the time of the year the data were acquired. We presume that the depth differences between different acquisition tracks must be small and relatively smooth in time. That is why we propose introducing a leaky integration operator (Claerbout, 1992) to model these secular data variations within our inversion scheme.

Abandoned strategy for attenuating tracks

An earlier strategy to remove the ship tracks is to filter the residual during the inversion as follows (Fomel and Claerbout, 1995):

$$\begin{aligned} \mathbf{0} &\approx \mathbf{r}_d = \mathbf{W}\frac{d}{ds}(\mathbf{B}\mathbf{H}^{-1}\mathbf{p} - \mathbf{d}) \\ \mathbf{0} &\approx \epsilon \mathbf{r}_p = \epsilon \mathbf{p}, \end{aligned} \tag{8}$$

where d/ds is the derivative along the track. The derivative removes the drift from the field data yet preserves the geological features. One consequence of the derivative is that it creates more glitches and spiky noise at the track ends and at the bad data points. In addition, the derivative might induce a loss of resolution in the final image. Both effects are illustrated in Figure 6, where we display the estimated $\mathbf{p}$ after inversion with the fitting goals in equation (8). The tracks have been attenuated, as expected. However, important geological features in the middle of the lake and on the shore are lost. In addition, the map is more noisy because of the aggravating effect of the derivative on bad data points. We do not fully understand why this approach works badly. One possible explanation is that

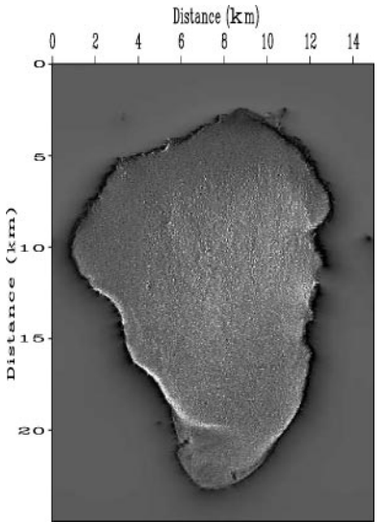

Figure 6. Estimated $\mathbf{p}$ after inversion with the fitting goals in equation (8). The derivative removes the tracks but creates a noisy image with a loss of resolution.

the conditioning of our problem with the operators $\mathbf{W}$, d/ds, $\mathbf{B}$, and $\mathbf{H}^{-1}$ worsens, making the optimization very difficult. Theoretically, we could estimate PEF from the data residual and use it as a data residual weight within the inversion. Unfortunately, the PEF would probably introduce glitches with good data because it would spread farther in space than the derivative.

Brown (2001) proposes estimating systematic errors between tracks by analyzing measurements at points where the acquisition swaths cross. This approach has the advantage of preserving the resolution of the depth map compared to the derivative along the tracks. Brown (2001) uses this idea as a preprocessing step, however. The comparison of Figure 4c and Figure 6 teaches us two lessons. First, the filtering approach with the derivative along the tracks does not produce a good image of the bathymetry (Fomel and Claerbout, 1995). Second, based on Brown's (2001) idea, we propose introducing an operator that will adaptively model and subtract the systematic shift within the inversion scheme. In the next section, we show that by incorporating a modeling operator for the drift in the data, we can effectively remove the ship tracks without any loss of resolution in the estimated depth map.

A new fitting goal

Now we present our idea of removing the tracks by adaptively subtracting them within our inversion scheme. Building on Nemeth et al. (2000), we introduce a modeling operator for the ship tracks inside our fitting goal in equation (6) as follows:

$$\begin{aligned} \mathbf{0} &\approx \mathbf{r}_d = \mathbf{W}(\mathbf{B}\mathbf{H}^{-1}\mathbf{p} + \lambda \mathbf{L}\mathbf{q} - \mathbf{d}), \\ \mathbf{0} &\approx \epsilon_1 \mathbf{r}_p = \epsilon_1 \mathbf{p}, \\ \mathbf{0} &\approx \epsilon_2 \mathbf{r}_q = \epsilon_2 \mathbf{q}, \end{aligned} \tag{9}$$

where $\mathbf{L}$ is a drift modeling operator (leaky integration), $\mathbf{q}$ is a new variable of our inversion, and λ is a balancing constant between gridding and noise modeling. We then minimize the misfit function,

$$g_2(\mathbf{p}, \mathbf{q}) = \|\mathbf{r}_d\|^2 + \epsilon_1^2 \|\mathbf{r}_p\|^2 + \epsilon_2^2 \|\mathbf{r}_q\|^2, \tag{10}$$

where $\mathbf{h} = \mathbf{H}^{-1}\mathbf{p}$ estimates the interpolated map of the lake. Again, we set $\epsilon_1 = \epsilon_2 = 0$, and we do not iterate to completion. For the operator $\mathbf{L}$, we choose a leaky integration operator such that $\mathbf{y} = \lambda \mathbf{L}\mathbf{q}$ is the portion of data value $\mathbf{d}$ that results from drift. Consistent with the way we use a rough variable $\mathbf{p}$ to represent the smooth water depth $\mathbf{h}$, we now represent (for the purpose of speeding iteration) $\mathbf{y}$ by a rougher function $\mathbf{q}$. The operator $\mathbf{L}$ has the following recursive form:

$$y_s = \rho y_{s-1} + q_s \quad s \text{ increasing along the data track.} \tag{11}$$

The parameter ρ controls the decay of the integration. For $\rho = 1$, leaky integration represents causal integration. The operator $\mathbf{L}$ is then appropriate to model the secular variations implied by the different seasonal and human conditions during data acquisition. We must choose a value of ρ that best represents the variations between the different tracks. This task is rather difficult to achieve: if ρ is too small, we might not be able to remove the drift; if ρ is too big, we might remove the drift and the bathymetry. Therefore, we carefully selected ρ by

starting from $\rho = 0.999$, interpolating with this value, looking at the final result, and decreasing ρ by 0.001 if necessary. We repeated this process until all tracks were attenuated. At the end of this exhaustive search, we found that $\rho = 0.99$ removes the tracks while preserving the bathymetry. We kept this value of $\rho = 0.99$ for all remaining results involving track attenuation and show that $\mathbf{L}$ removes most of the vessel tracks present in Figure 4.

The choice of λ in equation (9) is also critical. We tried different values by starting from a very small number and increasing it slowly. We then chose the smallest value that removed enough tracks in the final image ($\lambda = 0.08$). Nemeth et al. (2000) demonstrate that the noise (the tracks) and signal (the depth) can be separated in equation (9) if $\mathbf{L}$ and $\mathbf{BH^{-1}}$ do not model similar components of the data space. The parameter λ helps us to mitigate the possible crosstalk. A similar approach is used by Guitton (2002) to successfully remove ground roll on common midpoint gathers.

Figure 7 compares the estimated $\mathbf{p}$ with or without the attenuation of the vessel tracks. Figure 7b is essentially track free without loss of details, compared with Figure 7a. The difference plot in Figure 7c corroborates this and does not show any geological feature. We show in Figure 7b a linear feature that remains after we remove most of the spikes and the tracks. This event probably corresponds to a streak of bad data points that are incorrectly attenuated by our inversion schemes.

Comparing Figure 6 and Figure 7b, we see that the drift-modeling strategy [equation (9)] works much better than the noise-filtering strategy [equation (8)]. One possible explanation for the difference between the two results is that our modeling approach is more adaptive than the filtering of the residual. Indeed, by introducing the modeling operator, we basically look for the best $\mathbf{q}$ that models the drift of the data on each track at each point. The price to pay is an increase of the number of unknowns in equation (9). The reward is a surgically removed acquisition footprint. Notice that we can identify the ancient shorelines in the western and eastern parts of the lake very well.

To better understand what we are doing, we show in Figures 8 and 9 some segments of the input data $\mathbf{d}$, the estimated noise-free data $\mathbf{BH^{-1}p}$, the estimated secular variations $\lambda\mathbf{Lq}$ and the residual $\mathbf{BH^{-1}p} + \lambda\mathbf{Lq} - \mathbf{d}$ after inversion. The estimated noise-free data in Figures 8b and 9b show no remaining spikes. The effect of the track attenuation is more difficult to see because the amplitude of the drift is much smaller than the amplitude of the measurements. We can notice in Figure 8c that the estimated drift seems to have reasonable amplitudes: the average drift is around 15 cm for an accuracy of about 10 cm for the measurements, which is satisfactory. Also, the estimated drift is relatively constant throughout Figure 8c. Now, if we look at the estimated drift for another portion of the data (Figure 9c), we notice the drift has more variance than in Figure 8c and oscillates between 0 and 2 m, which is a lot. In addition, the estimated drift seems to follow the bathymetry of the lake in Figure 9a.

Looking closely at the residual (Figure 9d), we notice the drift is large where the data are noisy (Figure 9a). It is possible that the day of acquisition was very windy, a common weather condition for the Sea of Galilee (Volohonsky et al., 1983). Thus, the wind forces the water to pile up on one side of the lake, which can explain the lower water level on the other side. In addition, the strong wind in the middle of the lake induces noisy measurements because of the waves and the ship's erratic movement. It is also possible that the depth sounder was not working properly that day and had problems correctly measuring the deepest part of the lake. These causes could probably explain the shape and amplitude of the estimated drift in Figure 9c, but we cannot be absolutely sure. Unfortunately, no daily logs of the survey were kept to better interpret our results, especially for such a noisy data set.

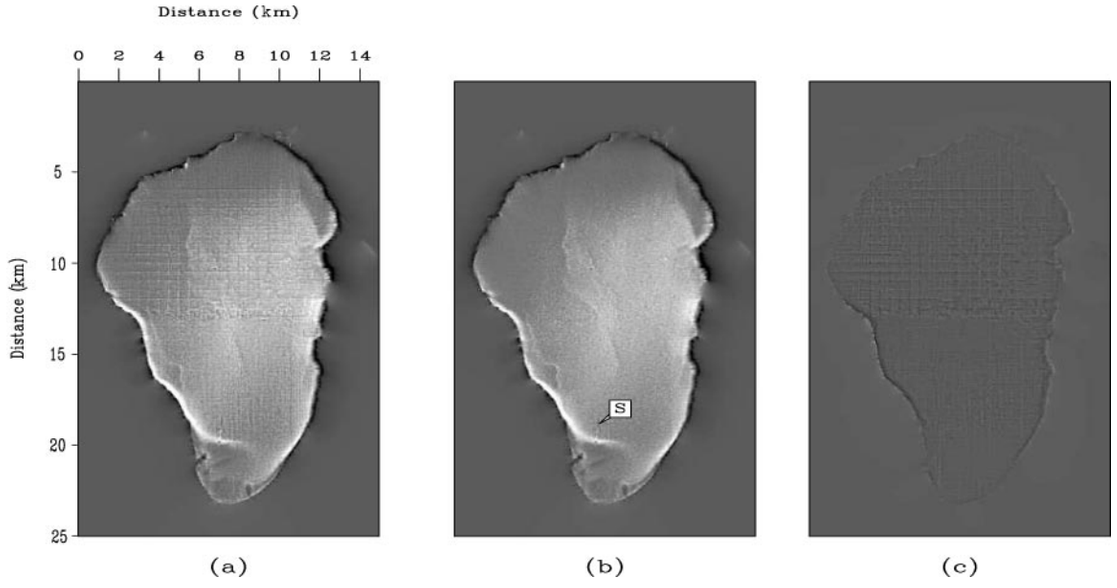

Figure 7. (a) Estimated $\mathbf{p}$ (a) without [i.e., equation (6)] and (b) with attenuation [i.e., equation (9)] of the tracks. S points to a linear, noisy feature that is not removed by the ℓ^1 norm or the modeling of the tracks. (c) Difference between (a) and (b); no geological feature has leaked in.

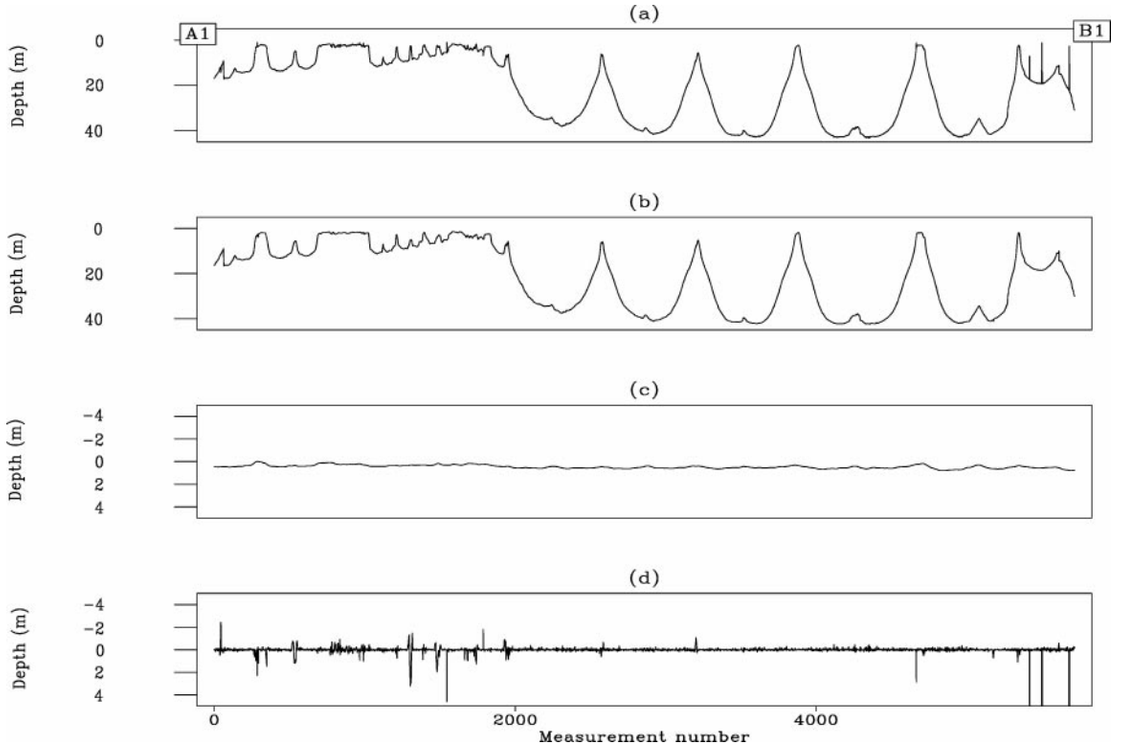

Figure 8. (a) Input data acquired between $A1$ and $B1$ in Figure 2. The ship is moving approximately bottom to top, going eastward from $A1$ to $B1$. (b) Estimated $\mathbf{BH}^{-1}\mathbf{p}$ after inversion, i.e., the estimated noise-free data. (c) Estimated drift after inversion. (d) Data residual after inversion. The horizontal axis represents the measurement number.

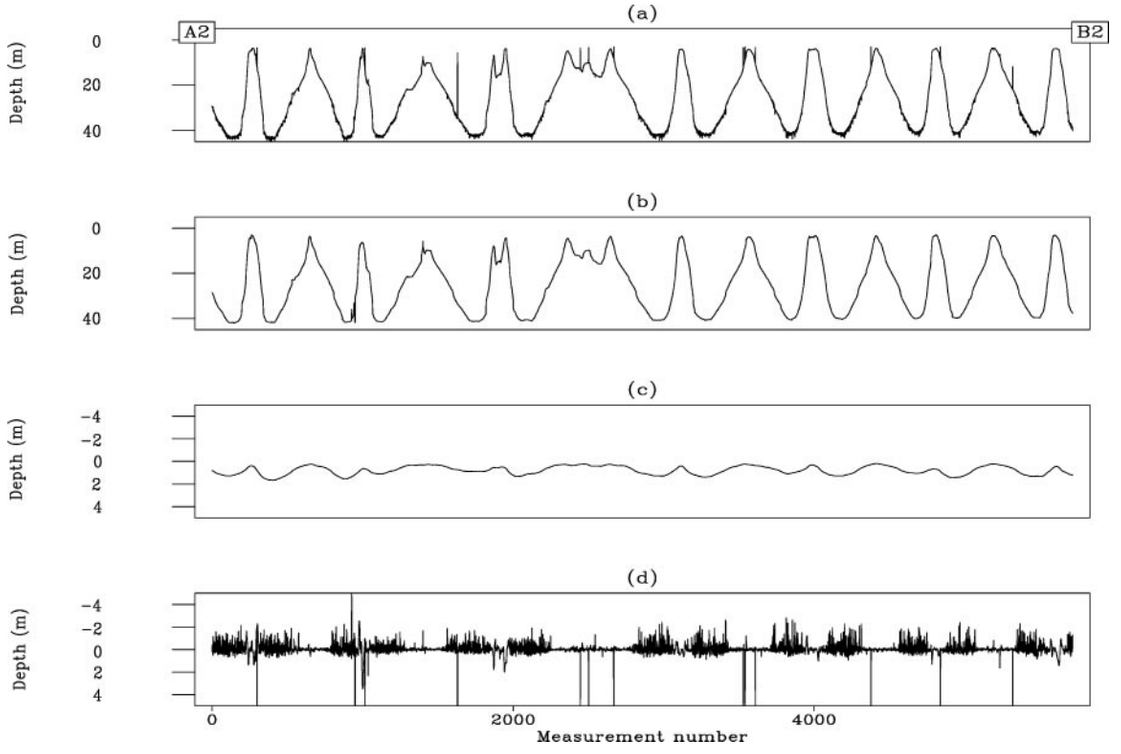

Figure 9. (a) Input data acquired between $A2$ and $B2$ in Figure 2. The ship is moving approximately right to left, going southward from $A2$ to $B2$. (b) Estimated $\mathbf{BH}^{-1}\mathbf{p}$ after inversion, i.e., the estimated noise-free data. (c) Estimated drift after inversion. (d) Data residual after inversion. The horizontal axis represents the measurement number.

We have shown that we can effectively subtract the tracks without any loss of resolution by introducing a noise modeling operator within our inversion scheme. In the next section, we go one step further and attempt to improve our result by using a PEF instead of the helix derivative as a preconditioner in equation (9). We show that the PEF helps to improve the contrast in some locations of the lake while decreasing the frequency content.

APPROXIMATING THE MODEL COVARIANCE WITH A PEF

The ideal regularization operator (squared) is the model covariance matrix (Tarantola, 1987). Estimating this matrix is not straightforward. We often approximate the model covariance matrix with roughening operators like the Laplacian or the derivative. Claerbout and Fomel (2002) advocate that, in principle, an ideal regularization operator is a PEF estimated from an a priori model. In this section, we test the idea of using a PEF instead of the helix derivative for the regularization operator.

Starting from equation (9), we replace the helix derivative $\mathbf{H}$ with a PEF termed $\mathbf{A_m}$, estimated from a model of the lake bathymetry. Since we do not know a priori the exact $\mathbf{h}$, a PEF is computed from the depth map $\mathbf{h}$ estimated in equation (9) with the helix derivative. We estimate a 3×4 filter. This procedure can be interpreted as a bootstrapping of the model covariance matrix. Including $\mathbf{A_m}$ into equation (9), we have the new fitting goals:

$$\begin{aligned} \mathbf{0} &\approx \mathbf{r}_d = \mathbf{W}\left(\mathbf{B}\mathbf{A}_\mathbf{m}^{-1}\mathbf{p} + \lambda\mathbf{L}\mathbf{q} - \mathbf{d}\right), \\ \mathbf{0} &\approx \epsilon_1 \mathbf{r}_p = \epsilon_1 \mathbf{p}, \\ \mathbf{0} &\approx \epsilon_2 \mathbf{r}_q = \epsilon_2 \mathbf{q}. \end{aligned} \tag{12}$$

Again, we set $\epsilon_1 = \epsilon_2 = 0$ and estimate $\mathbf{p}$ and $\mathbf{q}$ with conjugate gradients.

In Figure 10b, we show the estimated $\mathbf{p}$ with the PEF as a preconditioning operator. To increase the contrast inside the lake, we apply a weighting function on Figure 10b that boosts the low values in the middle of the sea and deemphasizes the seashores. We tried a similar weighting function on the result with the helix derivative but with no improvement in the lake image. Figure 10c displays the final result after weighting Figure 10b. The contrast in Figure 10c increases compared with Figure 10a (shown as HC). The ridge in the middle of the lake is no better defined with the PEF than with the helix derivative. In general, the final image with the PEF has a lower frequency content than in Figure 10a. The PEF tends to have a broader impulse response than the helix derivative (shown as LF), thus leading to an image with a low-frequency content.

Therefore, the approximation of the inverse covariance operator with a PEF appears to be a good idea: it increases the contrast of some interesting features. Nonetheless, the final result with the PEF has a lower frequency content as well. One major shortcoming with this last approach is that we need a good a priori $\mathbf{h}$ to estimate $\mathbf{A_m}$. Estimating $\mathbf{A_m}$ from the depth map obtained with the helix derivative might not be ideal. In addition, because the the lake has lateral and vertical variations in its bathymetry, we should estimate many PEFs, not just one. All of these suggestions are possible areas of improvement with the Sea of Galilee data set.

CONCLUSION

The interpolation of the Sea of Galilee data set on a regular grid is addressed as a noise attenuation problem. The combination of preconditioning and IRLS with the proper weighting function greatly reduces glitches and spikes. In addition, the introduction of a noise modeling operator that accounts for the inconsistency between depth measurements on different tracks greatly reduces the acquisition footprint. These two strategies

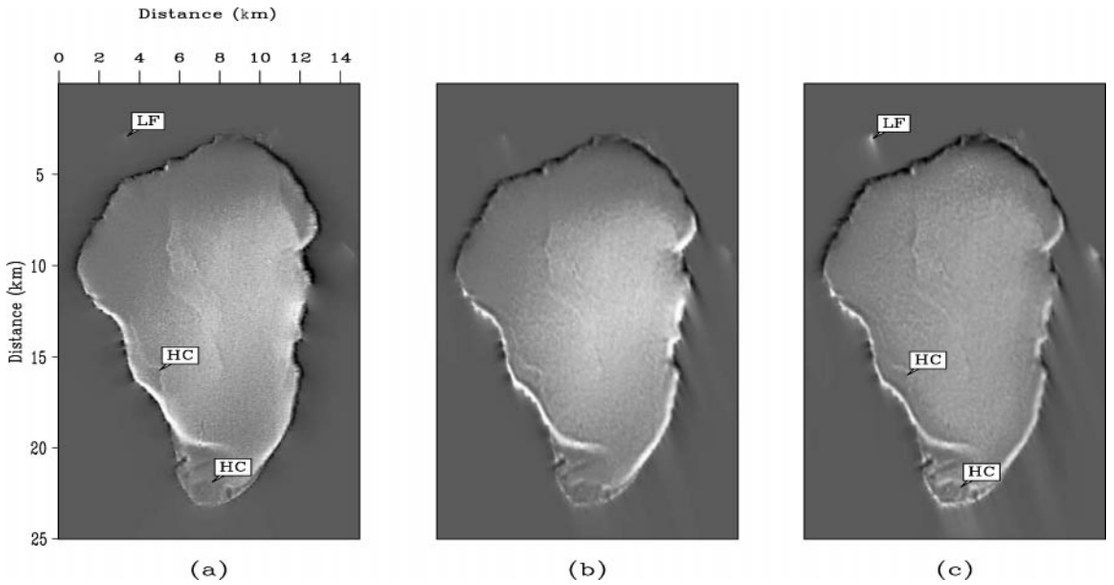

Figure 10. Estimated $\mathbf{p}$ with (a) the helix derivative (same as Figure 7b) and (b) the PEF. (c) Scaled version of (b) to increase the contrast inside the lake. The three plots are clipped to the 100th percentile. The term HC points to locations of higher contrast when the PEF is utilized for the inverse covariance operator. On the contrary, LF shows that the PEF slightly decreases the resolution (frequency content) of the final image by spreading information farther than the helix derivative.

help to unravel meaningful geological features inside the lake and ancient shorelines. Our last attempt for improving the final image with a PEF as a preconditioner produces a map with more contrast at the lakeshores but with a lower frequency content than the helix derivative. We recommend more work in that direction.

The lessons we learn from processing the Sea of Galilee data set can be reused in our daily geophysical work. We learn that the residual should always be looked at to derive the correct weighting functions. It also seems to teach us that it is better to model and subtract the noise than to try to filter or weight it out of the residual.

ACKNOWLEDGMENTS

We thank Zvi Ben-Avraham of Tel-Aviv University for providing the data set and the sponsors of the Stanford Exploration Project for supporting this research. We also thank Kurt Marfurt and the reviewers for their constructive remarks.

REFERENCES

Britze, P., 1998, Interpreter's corner—Seismic mapping using trend-form gridding: The Leading Edge, **17**, 606–608.
Brown, M., 2001, Estimation of systematic errors in tracked datasets using least squares crossing point analysis: SEP Annual Report, **110**, 123–132.
Bube, K. P., and Langan, R. T., 1997, Hybrid λ^1/λ^2 minimization with applications to tomography: Geophysics, **62**, 1183–1195.
Chapman, N. R., and Barrodale, I., 1983, Deconvolution of marine seismic data using the l1 norm: Geophysical Journal of the Royal Astronomical Society, **72**, 93–100.
Chemingui, N., and Biondi, B., 2002, Seismic data reconstruction by inversion to common offset: Geophysics, **67**, 1575–1585.
Claerbout, J. F., 1992, Earth sounding analysis, processing versus inversion: Blackwell Scientific Publications.
——— 1998, Multidimensional recursive filters via a helix: Geophysics, **63**, 1532–1541.
Claerbout, J. F., and Fomel, S., 2002, Image estimation by example: Geophysical soundings image construction (class notes): http://sepwww.stanford.edu/sep/prof/index.html.
Crawley, S., 2000, Seismic trace interpolation with nonstationary prediction-error filters: Ph.D. thesis, Stanford University.
Duijndam, A. J. W., Volker, A. W. F., and Zwartjes, P. M., 2000, Acquisition footprint reduction—Least squares migration versus reconstruction: 62nd Meeting, European Association of Geoscientists and Engineers, Extended Abstracts, session P0128.
Fomel, S., 2001, Three-dimensional seismic data regularization: Ph.D. thesis, Stanford University.
Fomel, S., and Claerbout, J. F., 1995, Searching the Sea of Galilee: the splendors and miseries of iteratively reweighted least squares: SEP Annual Report, **84**, 259–270.
Guitton, A., 2002, Coherent noise attenuation using inverse problems and prediction-error filters: First Break, **20**, 161–167.
Guitton, A., and Symes, W., 2003, Robust inversion of seismic data using the Huber norm: Geophysics, **68**, 1310–1319.
Guspi, F., and Introcaso, B., 2000, A sparse spectrum technique for gridding and separating potential field anomalies: Geophysics, **65**, 1154–1161.
Harlan, W. S., 1995, Regularization by model parameterization: http://www.billharalan.com/pub/papers/regularization.pdf.
Karpushin, A., and Brown, M., 2001, Whitening track residuals with PEFs in IRLS approach: SEP Annual Report, **110**, 133–140.
Nemeth, T., Sun, H., and Schuster, G. T., 2000, Separation of signal and coherent noise by migration filtering: Geophysics, **65**, 574–583.
Rickett, J., Guitton, A., and Gratwick, D., 2001, Adaptive multiple subtraction with non-stationary helical shaping filters: 63rd Meeting, European Association of Geoscientists and Engineers, Extended Abstracts, session P167.
Schuster, G. T., and Liu, Z., 2001, Seismic array theorem and rapid calculation and acquisition footprint noise: Geophysics, **66**, 1843–1849.
Tarantola, A., 1987, Inverse problem theory: Elsevier.
Trad, D., Ulrych, T., and Sacchi, M., 2003, Latest views of the sparse Radon transform: Geophysics, **68**, 689–399.
Volohonsky, H., Kaplunovsky, A., and Serruya, S., 1983, Storms on Lake Kinneret: Observations and mathematical model: Ecological Modelling, **18**, 141–153.