Stanford Exploration Project, Report 94, November 11, 1997, pages 189–205

Short Note

Pre-whitening and coherency filtering

Matthias Schwab[1]

keywords: *prediction error filter, deconvolution, coherency attribute*

INTRODUCTION

Ultimately, a seismic image serves interpreters as a means of building a geological model of the subsurface. Automatic edge detection schemes can help to produce images that emphasize critical geological features such as faults and river channels.

Showing some striking coherency images of complex geological structures, Bahorich and Farmer (1995) brought the coherency attribute to the attention of the geophysical community at large. The initial method scanned dips within a local image volume, estimated the dominant planar contribution and subtracted it. Gersztenkorn (Gersztenkorn and Marfurt, 1996) improved the original method by using the eigenvector representation of a local image volume to remove the coherent contribution. Luo et al. (1996) utilize two alternative shemes to yield similar coherency images. Luo's first scheme is based on simple, normalized trace differences while his second scheme calculates the spatial derivative of the traces' instantaneous phase.

In a previous article (Schwab et al., 1996), I used Prediction Error (PE) filtering to remove the coherent component of a local seismic image volume. PE filters are data adaptive and tend to whiten the spectrum of their input data (Claerbout, 1994). In another article in this report (Schwab, 1997), I discuss a set of PE filters that is designed to remove the plane layer component of a local data volume. Unfortunately, this set of PE filters removes almost all the coherent information and leaves only a white output image that is useless for interpretation.

In this paper, I describe my attempt to control the spectrum of the output of the PE filtering more carefully. In particular, I will attempt to remove horizontally coherent events while preserving most of the temporal coherency. The resulting images are rather disappointing. In general, pre-whitening in the time domain allows us to shape the temporal image spectrum. However, targets, such as fault reflections and salt dome flanks, were spatially predicted and removed along with the locally planar layered volumes. The final coherency images are too white to be useful.

[1]**email:** matt@sep.stanford.edu

In the first section, I will discuss how the subsequent application of two PE filtering steps permits us to shape the output spectrum. In the second section, I apply PE filtering to a few field data examples and discuss the individual results. PE filtering fails to produce the desired coherency images. In the conclusion, I will discuss my difficulties in finding an optimal parameter choice.

PREDICTION ERROR FILTERING

An interpreter desires a detailed image of the subsurface's features of interest, such as river channels or faults. Simultaneously, the interpreter may want to remove other features, such as a plane layer contribution.

A PE filter removes the predictable parts of an input image and creates a white output. Plane layer contributions are highly redundant and predictable and therefore likely to be removed by a PE filter. Unfortunately, even faults and river channels are predictable since they are continuous surfaces that locally are approximated by planes. I attempted to implement a PE filter scheme that is able to remove the plane layer volume, but does not remove planar discontinuities such as faults and river channel flanks.

In a more theoretical accompanying report (Schwab, 1997), I describe a set of three 2-D filters that are designed to remove plane layer volumes without removing planar discontinuities such as fault reflections. Additionally, the article suggests how to unite the individual three images of the three filters to a single final image volume.

The filter results in earlier experiments (Schwab et al., 1996) were disappointing since, despite their design, the filters removed all planar features that I attempted to isolate. In seismic sections, faults, flanks of river channels or salt domes are rarely pure planar features. Their limited extension out of a plane are sufficient for the PE filters to predict and remove the target events. Consequently, I decided to restrict the filter even more.

I know three methods to restrict a filter's ability to remove predictable input components:

- Deconvolve first and then apply the desired spectrum
- Restrict the PE filter's shape (such as shortening it or introducing a gap of zeroes between the first nonzero coefficient and the variable coefficients).
- Pre-filter the input before finding the PE filter

In my earlier coherency paper (1996), I experimented with filter gaps and filter shape (second approach). Since these approaches failed to restrict the filter sufficiently, I attempt in this paper to pre-whiten the filter input (third approach).

THE THEORY OF PRE-WHITENING

The central idea of pre-whitening is to use two PE filters. Each PE filter minimizes the mean square energy of its output given its input. However, the second filter is given the output of the first filter. Once we have found the second filter, we apply it to the original input (the input that we used to train the first filter). Obviously, the second filter will not be optimal for removing the predictable events of the original input. As a matter of fact, the second filter is blind with regard to the events the first filter removed from the data. The second filter applied to the original input will remove the components it was trained to remove, but it will preserve the components that the first filter is capable of removing.

Imagine that an input time series contains three predictable components. Given an input time series and a given filter length, we find the first optimal PE filter. The data component this first filter removes is the first predictable component.

If the spectrum of the input time series is Y_1, then a PE filter P_1 corresponds to

$$P_1 = \frac{1}{(< \bar{Y}_1 Y_1 >)^{1/2}}$$

where $(< \bar{Y} Y >)^{1/2}$ represents a smooth version of Y's spectrum. The output spectrum after such a PE filter step is

$$\begin{aligned} P_1 Y_1 &= \frac{Y_1}{(< \bar{Y}_1 Y_1 >)^{1/2}} \\ &= W_1. \end{aligned}$$

where W_1 represents a spectrum that tends to be white. If P_1 were a perfect PE filter, then W_1 would be white. But since our PE filter implementations are limited to a few coefficients, W_1 only tends to be white.

I define Y_2 oas the utput of the first PE filter

$$\begin{aligned} Y_2 &= \frac{1}{(< \bar{Y}_1 Y_1 >)^{1/2}} Y_1 \\ &= W_1. \end{aligned} \tag{1}$$

Consequently, the corresponding PE filter is

$$P_2 = \frac{1}{(\ll \bar{Y}_2 Y_2 \gg)^{1/2}}$$

where $\ll\gg$ represents a smoothing that might be different from the P_1 smoothing $<>$. Applying P_2 to Y_2 yields

$$\begin{aligned} P_2 Y_2 &= \frac{Y_2}{(\ll \bar{Y}_2 Y_2 \gg)^{1/2}} \\ &= W_2 \end{aligned}$$

where W_2 represents a spectrum that again only tends to be white. Substitution of the the expression 1 for Y_2 yields

$$W_2 = P_2 Y_2 = P_2 W_1$$

which indicates that W_2 is whiter than W_1.

If we apply the second PE filter, P_2 to the original data spectrum Y_1, we finally get

$$P_2 Y_1 = \frac{Y_1}{(\ll \bar{Y}_2 Y_2 \gg)^{1/2}}$$

The simple identity

$$Y_1 = \frac{Y_1}{(< \bar{Y}_1 Y_1 >)^{1/2}} (< \bar{Y}_1 Y_1 >)^{1/2}$$

allows us to introduce Y_2 according to expression 1

$$\begin{aligned} P_2 Y_1 &= P_2 \frac{1}{(< \bar{Y}_1 Y_1 >)^{1/2}} Y_1 (< \bar{Y}_1 Y_1 >)^{1/2} \\ &= P_2 Y_2 (< \bar{Y}_1 Y_1 >)^{1/2} \\ &= W_2 (< \bar{Y}_1 Y_1 >)^{1/2} \end{aligned}$$

Consequently, the resulting spectrum of the two step filter operation is the product of the first predictable spectral component $(< \bar{Y}_1 Y_1 >)^{1/2}$ and the rather white spectrum W_2. The second predictable component $(\ll \bar{Y} Y \gg)^{1/2}$ is removed. In that sense the cascading of PE filters implements a high resolution, data adaptive notch filter. The component removed, however, is not a fixed frequency component but a data dependent band of *predictable components*.

IMPLEMENTATION

Seismic image volumes are non-stationary. To estimate reasonable PE filters, I break the input image into small local volumes that are approximately stationary. After computing a coherency image using the method described above, I merge the local volumes to a single coherency image. Claerbout (Claerbout, 1994) implemented a method to split and merge a data volume into small local image volumes (patches).

To process an image volume, I split it into small local volumes, *subvolumes*. I then filter each subvolume individually and merge the resulting subvolumes back into a single output image volume. The split into individual subvolumes ensures that the local signal is approximately stationary and the filter step is optimal.

Figure 1 illustrates the data flow. The input data at the upper left hand is split into individual subvolumes (shaded rectangle). Next we find an optimal 1-D filter for each subvolume (symbolized by the elliptic symbol). The filtered subvolume is merged to the intermediate pre-whitened image volume at the top center of the flow chart.

Next we split the pre-whitened image into subvolumes (not necessarily of the same size as in the first filter step). We find the set of three 2-D filters as described in the companion paper (Schwab, 1997). The filter step yields three output images, which I discard.

The three PE filters are used to filter a corresponding subvolume of the original data, indicated by the elliptical operation symbol in the lower right corner of the diagram. This final filter step yields a single image output volume (shown at the bottom right) since it comprises a forward and adjoint filtering step.

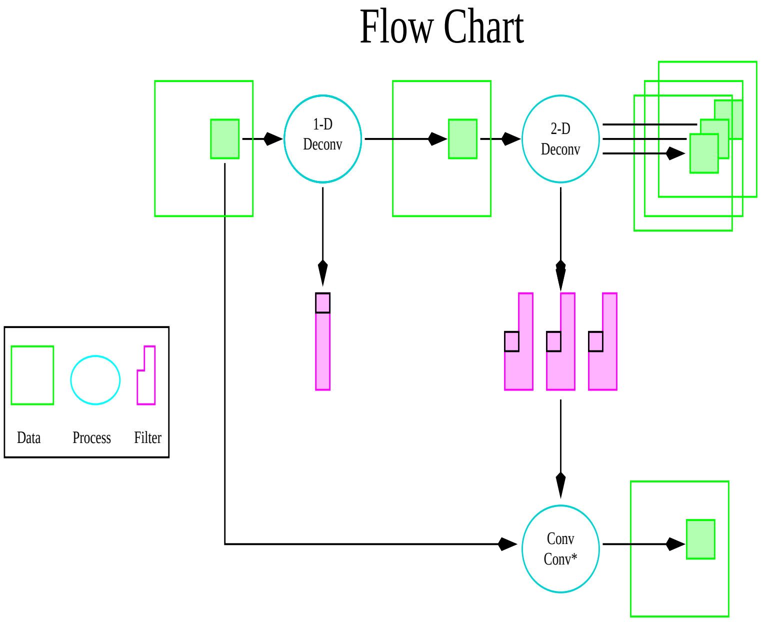

Figure 1: The diagram displays the processing flow to compute a coherency attribute. The process uses a two step PEF approach: the first PEF filter hides events from the second PEF filter. The second PEF filter ultimately applied to the original data will only remove predictable events that existed in the pre-filtered data set it was trained on. $\boxed{\text{matt1-flow}}$ [NR]

PARAMETER CHOICES

A user has the following options to increase the whiteness of a coherency image:

- decrease the size of the pre-whitening PE filter
- increase the size of the subvolume that we use to find the pre-whitening filter
- decrease the number of iterations for the pre-whitening PE filter
- increase the size of the final PE filter set
- decrease the size of the subvolume that we use to find the final filter set

- increase the number of iterations for the final PE filter

Several of these options contain several individual parameters, e.g., the subvolume size is defined by 3 integer values that define the subvolume's number of samples along each axis. The user has to set about a dozen parameters. Some of these parameters have obvious limits.

I chose the pre-whitening PE filter to be a 1-D deconvolution filter in time. This should yield optimal spatial whitening, while preserving a coherent waveform along the time axis.

The second PE filter step involves three 2-D PE filters (one in each Cartesian plane). I choose these filters to be two columns wide, which limits them to the removal of a single planar event. I choose the number of rows large enough to be able to predict the data's steepest dipping feature.

PE filter coefficients are found from the data by iteratively solving a minimization problem. An increase in the number of iterations usually improves the filter's ability to predict the input data. Since the filters are usually small, I decided to use a sufficient number of iterations to find the optimal least square filter. However, limiting the number of iterations is a convenient way to decrease a filter's ability to whiten the output.

An individual subvolume has to be large enough to determine the unknown PE filter coefficients by minimizing the convolution output. The subvolume serves as an expectation volume of a locally stationary signal.

Despite these obvious limits to the individual parameters, a twelve dimensional parameter space is very hard to explore exhaustively. The parameters are probably interdependent. For example, decreasing the individual subvolume size has an effect similar to increasing the number of iterations for finding the filter; both yield a more accurate filter estimation and therefore a whiter filter output. I tested several sets of parameters in the subsequent field data examples and I did not succeed creating a satisfying coherency image. The images that I show in the next section are representative.

FIELD DATA RESULTS

The following figures show four data sets during the various stages of processing. In each section, the first image shows the migrated image, the second image shows the pre-whitened data, and the last image shows the final coherency image. Each image is represented by a time or depth slide on the left and an in-line and cross-line image on the right.

In all final images the inline and cross-line section show significant coherency along each individual trace. However, the images show hardly any coherency in the spatial direction. All targets of interest were removed.

River channel

The main target in this migrated time-slice contains a meandering, buried river. The data quality is poor. My version of this data set is stored in a strange floating point representation. I nevertheless like to include this figure since it includes a nice example of a meandering river and since I had some of my best results with this data set.

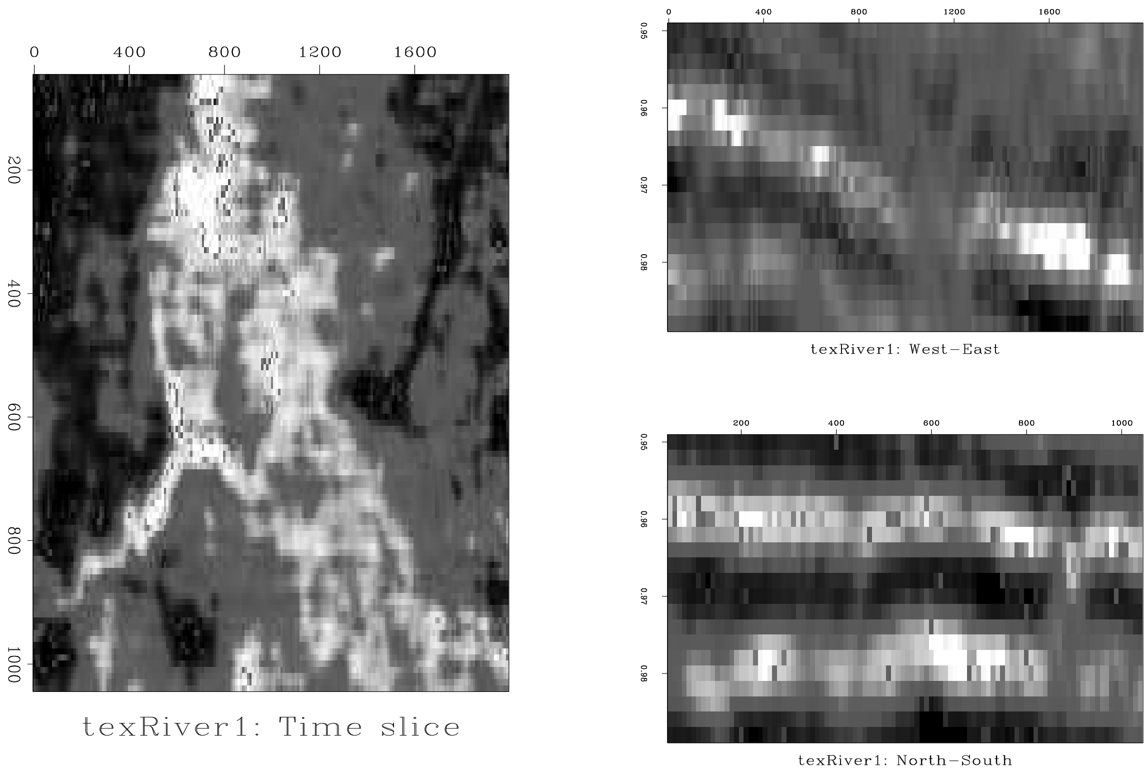

Figure 2: Migrated image. matt1-texRiverSect1 [CR]

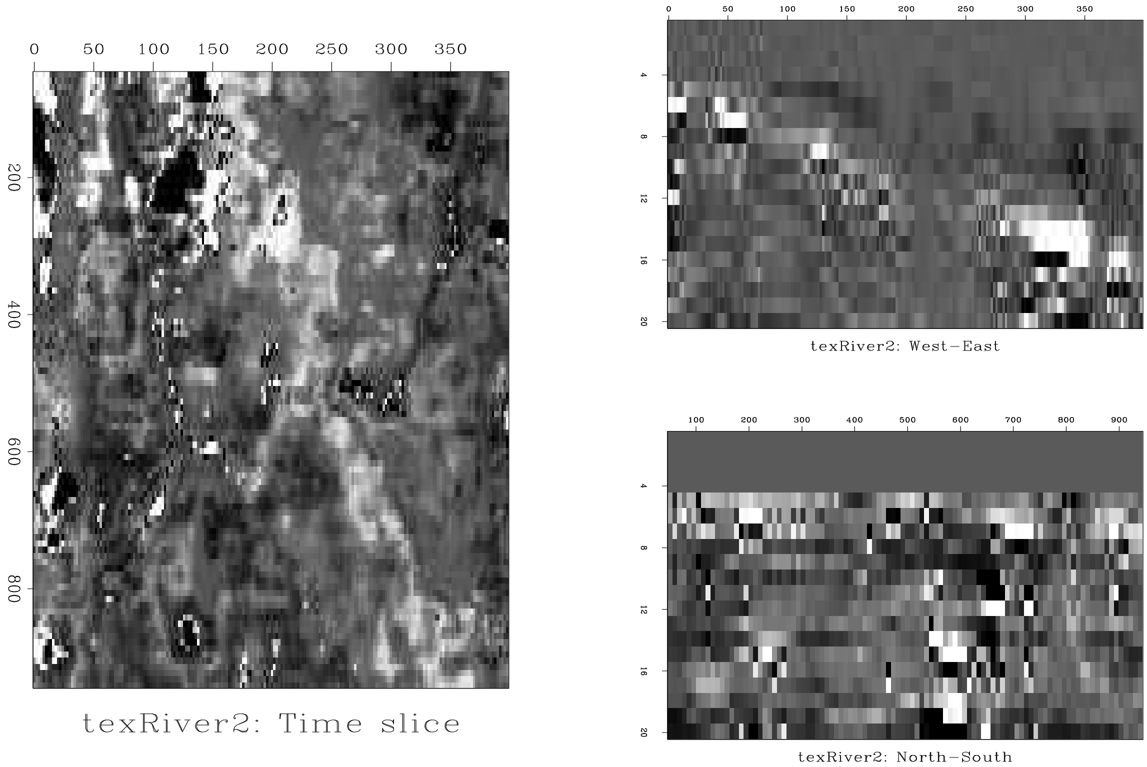

Figure 3: Pre-whitened image. matt1-texRiverSect2 [CR]

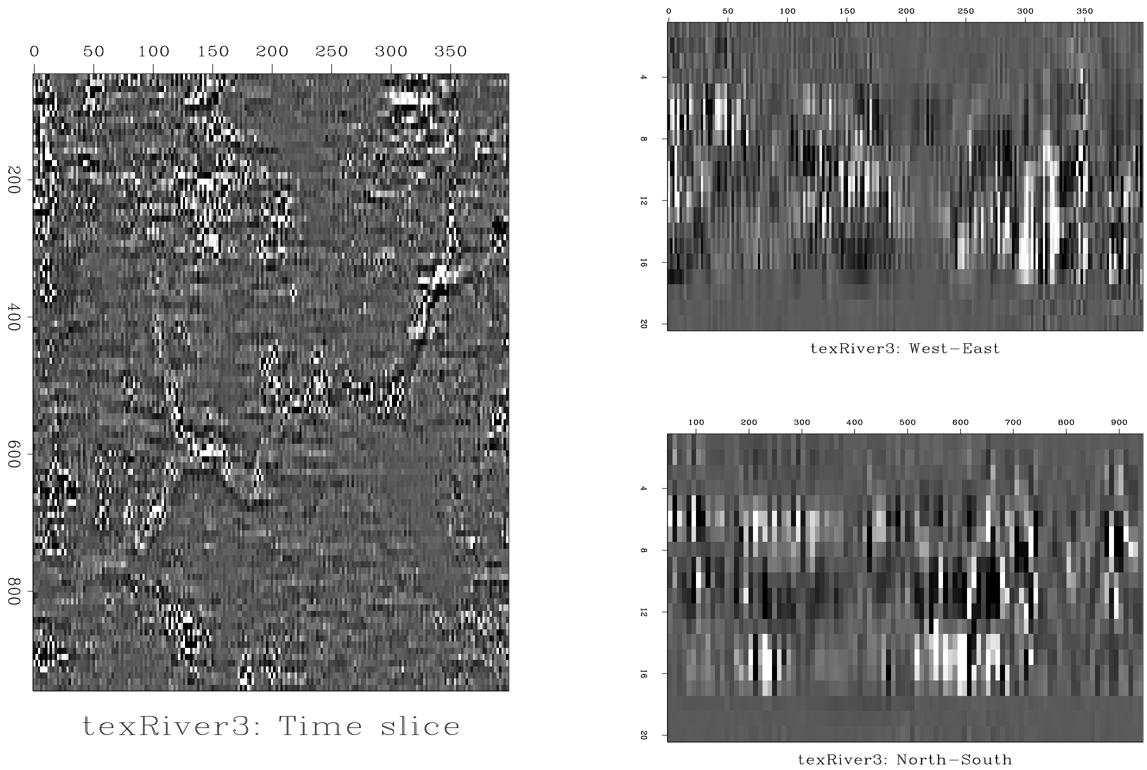

Figure 4: Coherency image. matt1-texRiverSect3 [CR]

Salt dome

The target here is a salt dome flank and a tangential fault. Unfortunately, the distance between depth slices is rather large and therefore the input image is hardly coherent in depth. The coherency computation is limited to depth slices, and so pre-whitening with a 1-D deconvolution in depth has little effect.

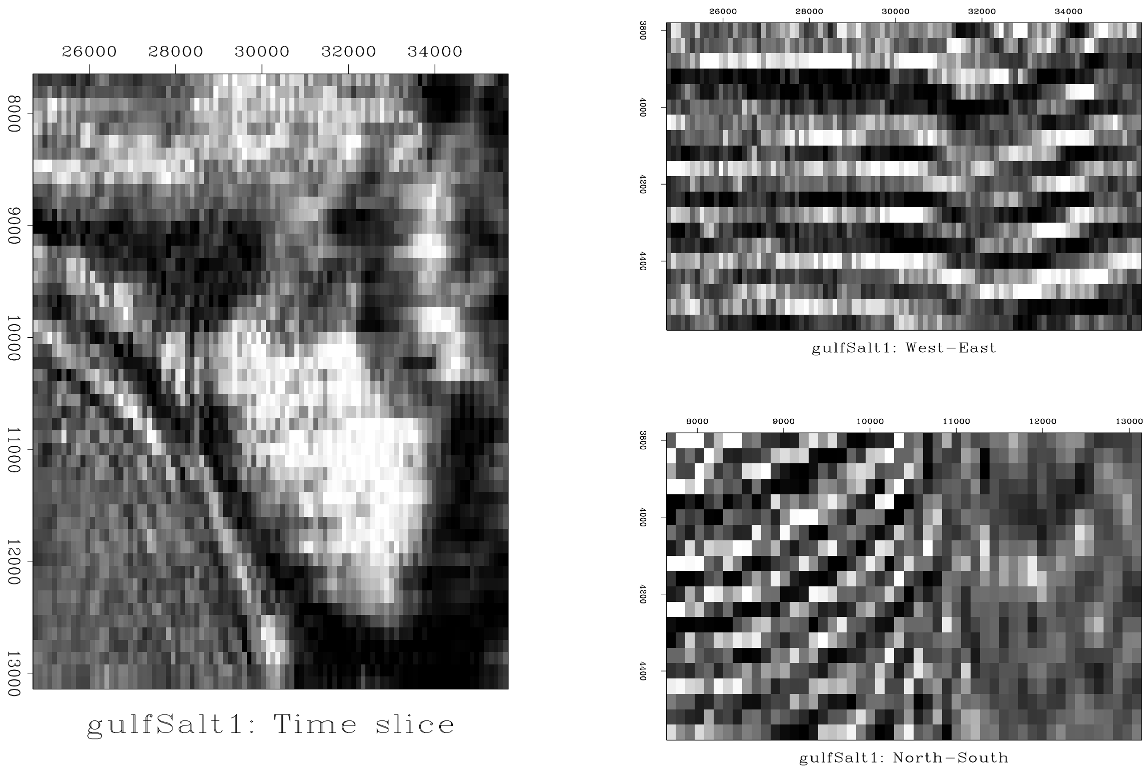

Figure 5: Migrated image. matt1-gulfSaltSect1 [CR]

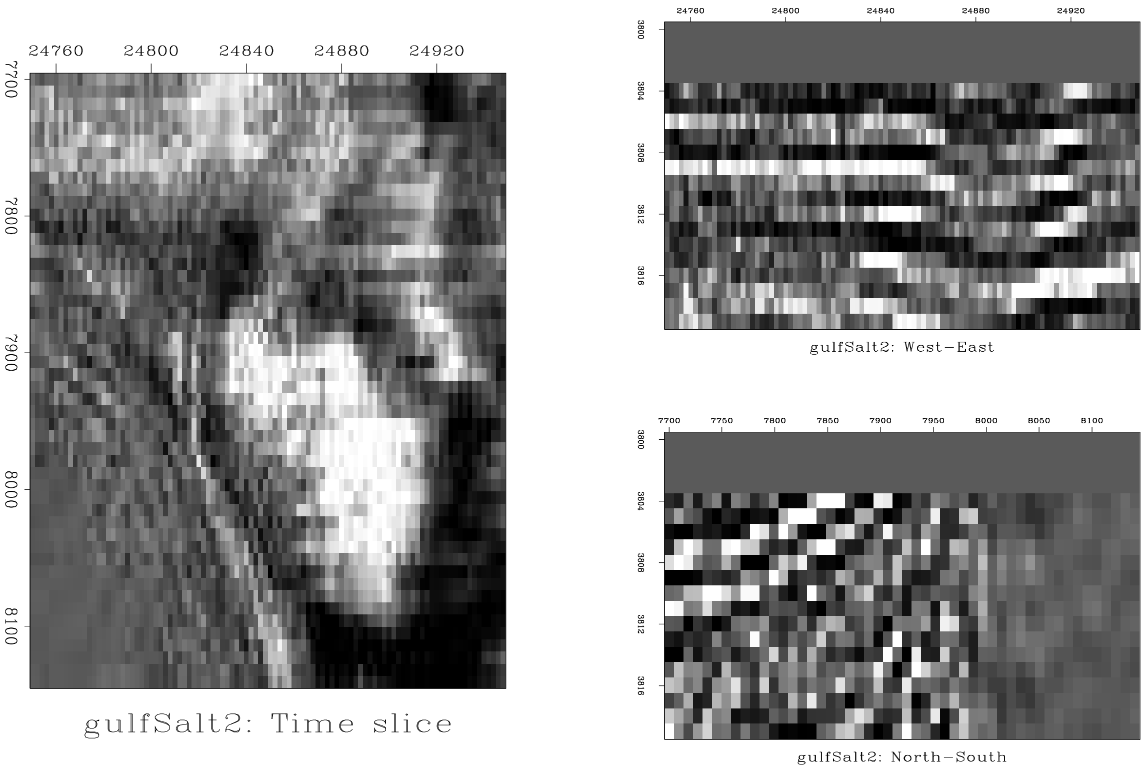

Figure 6: Pre-whitened image. $\boxed{\text{matt1-gulfSaltSect2}}$ [CR]

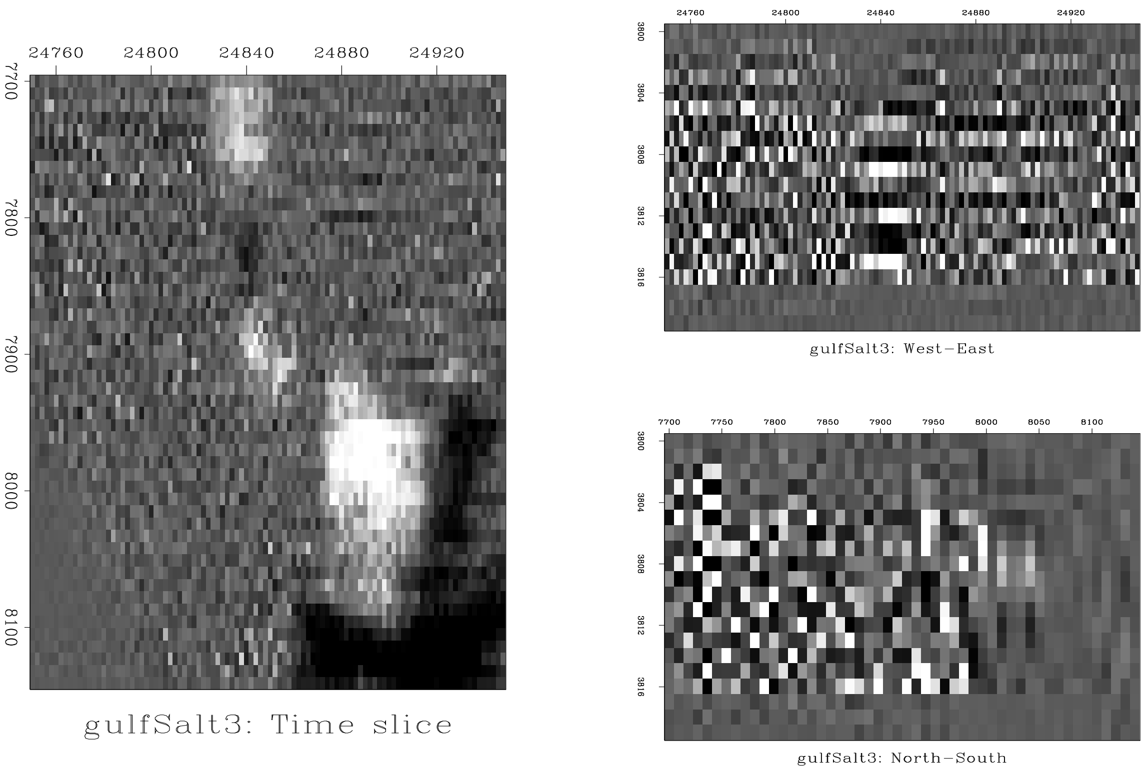

Figure 7: Coherency image. $\boxed{\text{matt1-gulfSaltSect3}}$ [CR]

Salt anticline

The target is a salt anticline with a normal fault at the anticline's tip. The anticline's flanks are rather steep and the fault zone appears to be rather predictable. The fault even shows some fault reflections.

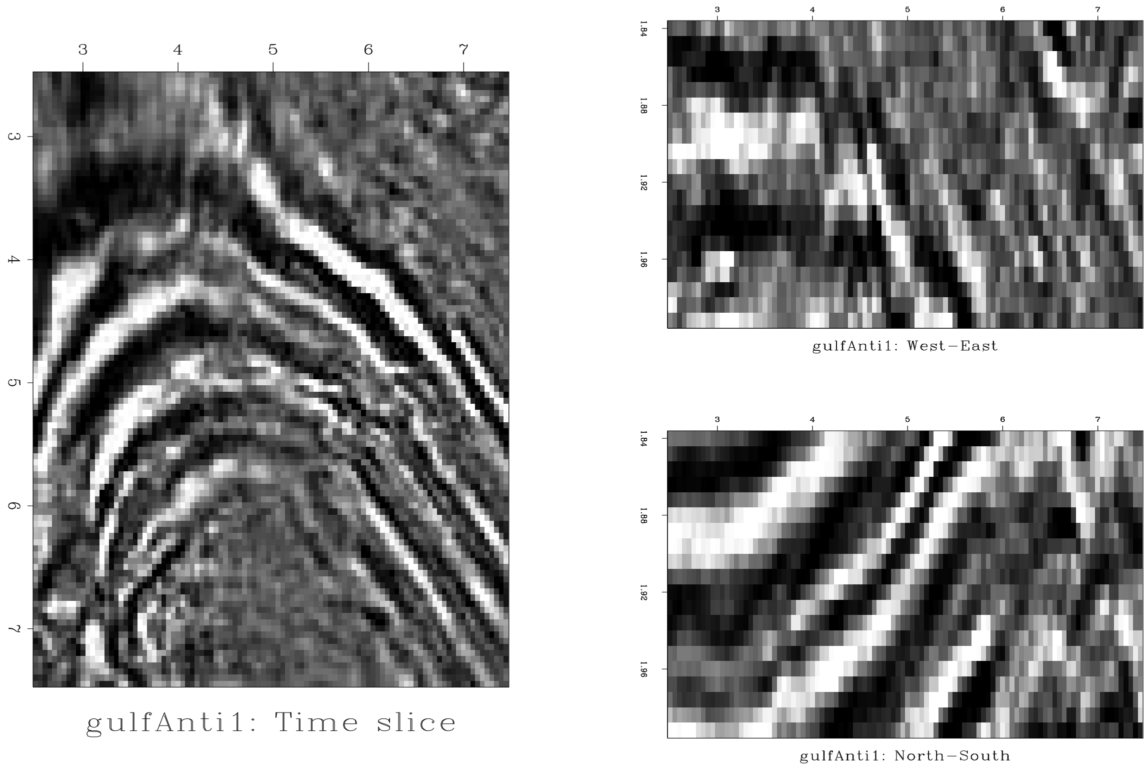

Figure 8: Migrated image. matt1-gulfAntiSect1 [CR]

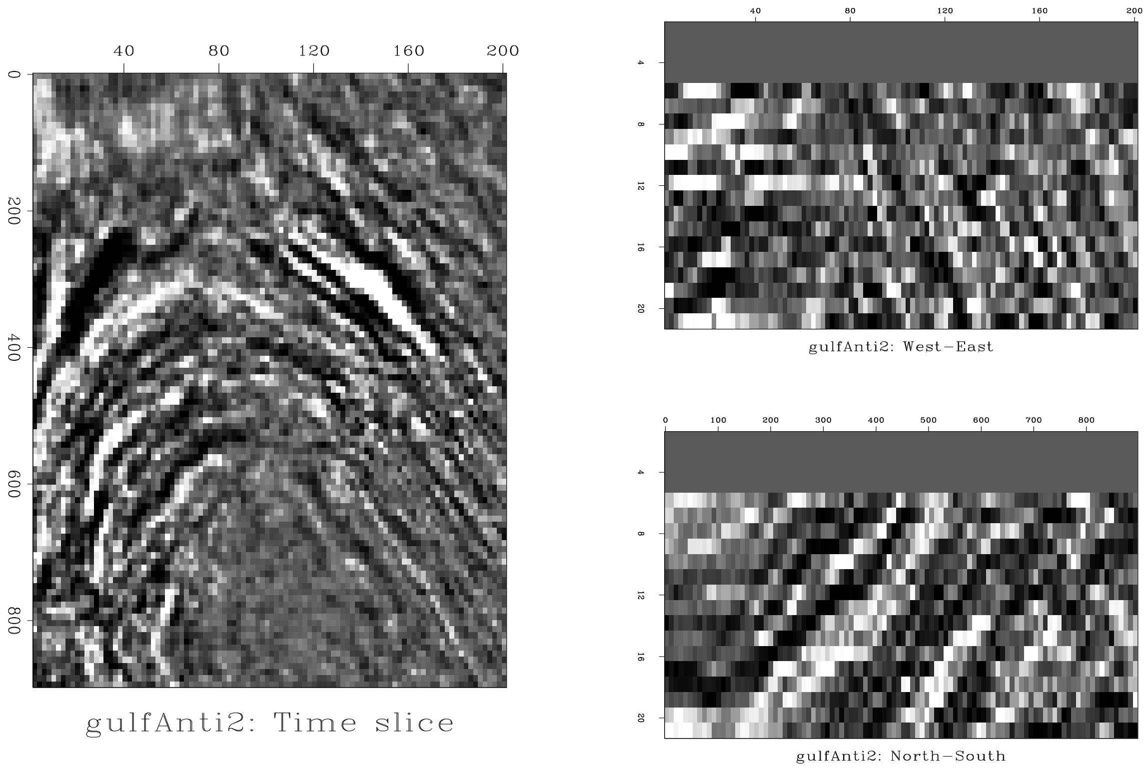

Figure 9: Pre-whitened image. matt1-gulfAntiSect2 [CR]

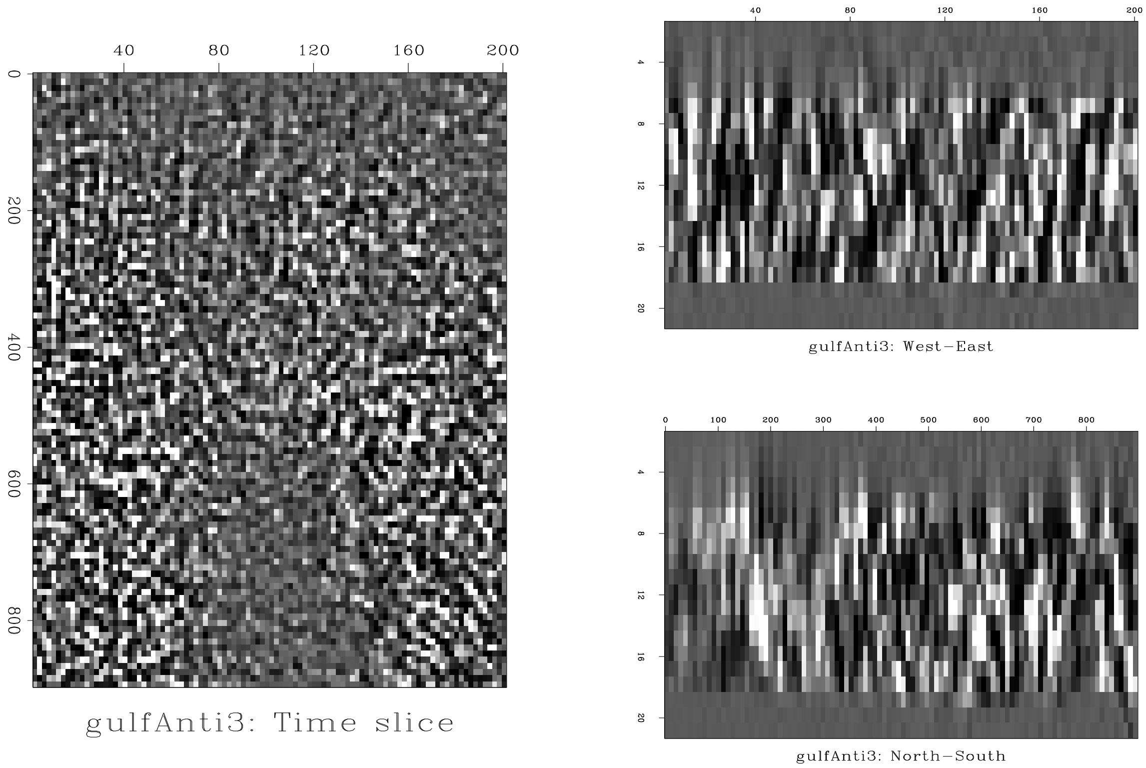

Figure 10: Coherency image. matt1-gulfAntiSect3 [CR]

Carbonate reefs

The target is a carbonate reef. The reflector slope steeply and the geology is complicated. I cannot find a particular target, but I can image that there are interesting features hidden in the jumble of continuous reflectors. Unfortunately, my coherency attribute computation predicts and removes all events.

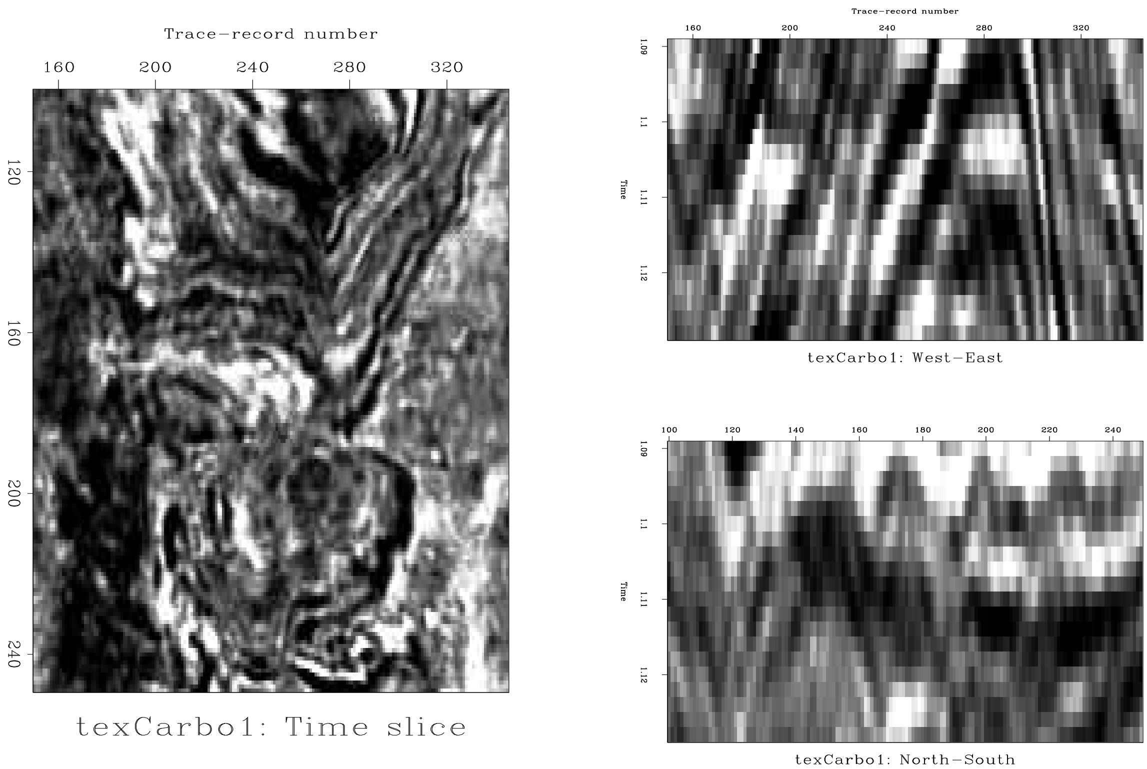

Figure 11: Migrated image. $\boxed{\text{matt1-texCarboSect1}}$ [CR]

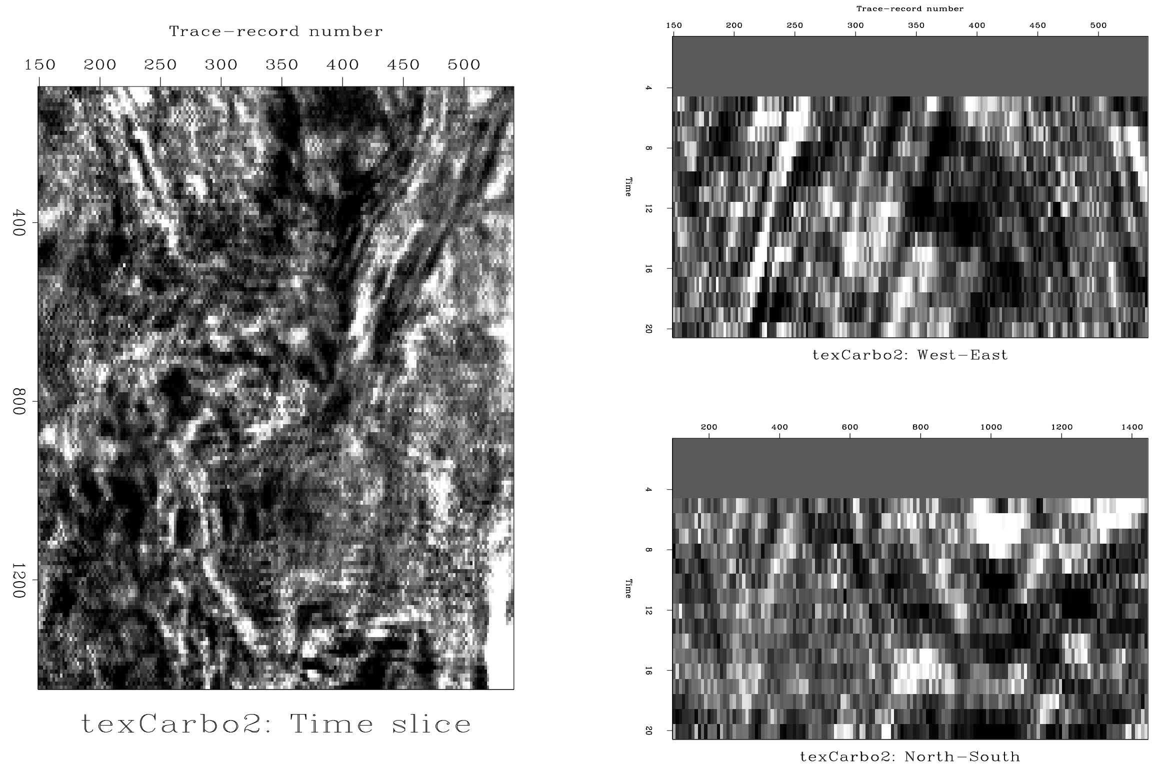

Figure 12: Pre-whitened image. $\boxed{\text{matt1-texCarboSect2}}$ [CR]

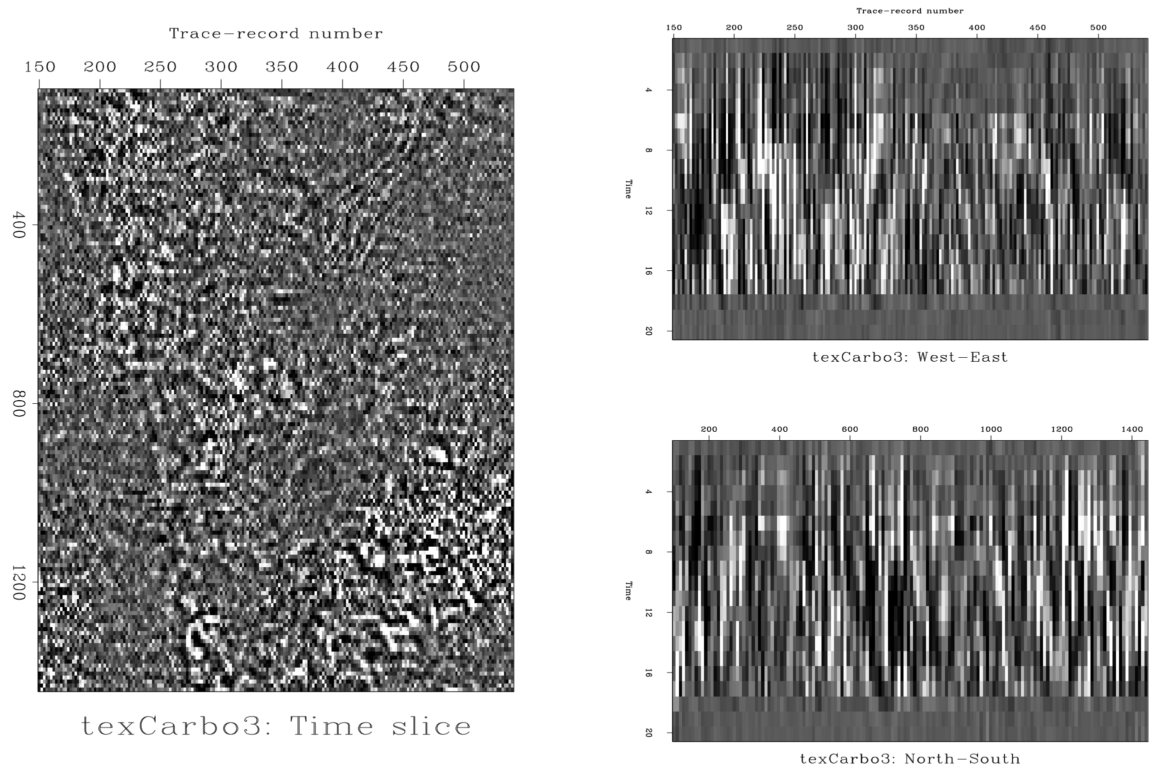

Figure 13: Coherency image. $\boxed{\text{matt1-texCarboSect3}}$ [CR]

CONCLUSION

An interpreter wants a 3-D image of the subsurface in which planar features (faults, salt dome flanks) are sharpened and coherent, while 3-D plane layer volumes are removed. My filter approach, however, removes all seismic events despite the fact that it is designed to remove only plane layers and to protect temporal coherency. Consequently, I failed to control the output spectrum sufficiently and the final images are so white that they cannot be used for geological interpretation.

As expected, temporal pre-whitening helped to preserve temporal coherency. However, the limited spatial coherency of *planar* seismic features enabled my PE filter to remove all features of interest.

The approach could possibly be enhanced by spatial pre-whitening or other methods of spectral spatial shaping (e.g. introduction of an additional filter gap). Unfortunately, such a spatial whitening further increases the number of parameters and complicates the pre-whitening filter step tremendously.

The number of interdependent parameters already inhibits an exhaustive exploration of all possible options. Having tested various sets of parameter combinations, I failed to produce attractive coherency images. Since the parameters are not independent of each other in their effect on the outcome, I doubt that a more exhaustive search would succeed.

Controlling an images output spectrum through PE filters has proven to be difficult. Better theoretical understanding of PE filters might clarify if a combination with a moving average approach would allow more finely tuned control.

Finally, I am considering building a series of small synthetic data volumes. Such synthetic image volumes would allow me to test the coherency attribute using PE filters or alternative methods in a more controlled setting.

REFERENCES

Bahorich, M., and Farmer, S., 1995, The coherency cube: The Leading Edge, **14**, 1053–1058.

Claerbout, J. F., 1994, Applications of Three-Dimensional Filtering: Stanford Exploration Project.

Gersztenkorn, A., and Marfurt, K., 1996, Eigenstructure based coherence computations: 66th Annual Meeting: Expanded Abstracts, Society of Exploration Geophysicists, 328–331.

Luo, Y., Higgs, W. G., and Kowalik, W. S., 1996, Edge detection and stratigraphic analysis using 3-D seismic data: 66th Annual Meeting: Expanded Abstracts, Society of Exploration Geophysicists, 324–327.

Schwab, M., Holden, C., and Claerbout, J., 1996, Revealing geological discontinuities by local plane wave suppression: SEP–**92**, 29–48.

Schwab, M., 1997, Cross product operator detects plane reflectors: SEP–**94**, 181–188.