Searching for similarity in a slab of seismic data

Stephen J. Hill, Colorado School of Mines, Golden, USA
Kurt Marfurt, University of Houston, USA
Satinder Chopra, Arcis Corporation, Calgary, Canada

In the last ten years, geoscientists developed many methods to interpret layered, 3D data. This article will describe some of the solutions, and a good place to begin is with a review of the challenges of 3D seismic interpretation.

Seismic interpreters describe the interior of a solid. That is a difficult goal for even an ordinary rock (Figure 1) that we can hold in our hands even though we can rotate it to gain an appreciation for its interior. We are even allowed to fracture the rock to investigate its interior. However, even with such advantages, we are challenged to recreate its detailed 3D structure in our minds.

Figure 1. *A Colorado rock. A U.S. penny is shown for scale.*

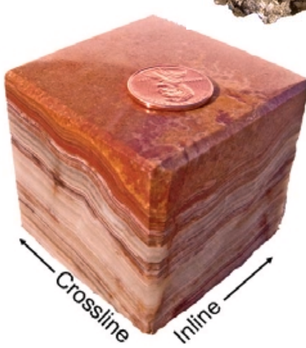

Figure 2. *Rock paperweight.*

A sedimentary rock paperweight (Figure 2) presents less of a challenge because it exhibits layering. This layering is critical for the application of our methods. We will use standard seismic terminology (inline, crossline, and depth) for its coordinates.

Figure 3 shows a depth slice (perpendicular to the depth axis) through this rock. In fact, Figure 3 is simply a photograph of the bottom of this rock paperweight. For this slice, the layering symmetry is not evident.

Figure 3. *Depth slice through paperweight.*

Figure 4 shows an inline slice of the rock. It is delightful that this inline slice of the paperweight looks so much like an inline slice through seismic data. The crossline slice reveals similar layering (Figure 5).

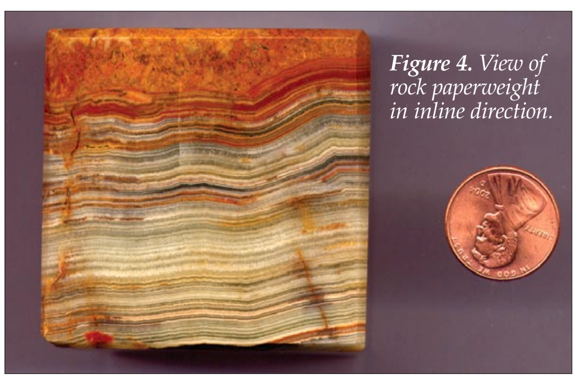

Figure 4. *View of rock paperweight in inline direction.*

During most of the history of seismic profiling, interpreters worked with 2D data that, at best, revealed discrete slices through the earth, such as shown in Figure 4 and Figure 5. From those views, seismic interpreters inferred the structure within the rock volume. Our paperweight shows faults, fractures, and changes of dip. But, when limited to these 2D views of the paperweight, the 3D picture of the actual locations of these geologic features exists only in our imagination.

The advent of 3D seismic put the entire cube at our disposal. We can use a workstation to slice and dice it to extract any view that tickles our fancy. For example, we can create horizontal slices, such as Figure 3, that, of course, we could not have obtained with 2D data. But, although the data and tools are better, the challenge to the interpreters remains: *How do we visualize, capture, and communicate a description of the interior of our rock paperweight, or of our seismic volume?* Surprisingly, real life in our 3D world does not give us many clues. Even though we occupy a 3D world, we see only 2D surfaces.

The first method of visualizing our 3D data extrapolates the procedures we used with 2D. We slice and dice the data, noting the depth (or traveltime) to a given horizon and, from that, produce a horizon map that represents a surface in the 3D volume. The arrow in Figure 6 points to just such a picked (green) horizon.

Capitalizing on lateral similarity. In Figure 6 we observe the paperweight's pervasive lateral similarity, the nearly horizontal layering. Nevertheless, our eyes and our interest are drawn to the departures from that lateral similarity. Even though this is just a paperweight, we start to ask: *Is that a fault? Is this a channel?* The departures from perfect lateral similarity tell us about the geology of this paperweight. That will be our motivation—to view changes in lateral similarity. Before presenting how that is done, we will focus more on "similarity."

We can map horizons because of the pervasive lateral similarity of the rocks and of the resultant seismic data. Figure 6 shows the direction of that similarity, with the vertical direction as the direction of dissimilarity. For layers close to the green (picked) horizon, the maximum direction of similarity is not along the horizontal inline direction but rather along the irregular green horizon. Therefore, we could first flatten to that green horizon in order to increase the degree of similarity in the inline direction to highlight departures

Editor's note: This article is related to the 2006 Distinguished Instructor Short Course "Seismic Attribute Mapping of Structure and Stratigraphy" which will be taught by Kurt Marfurt. SEG will also publish a reference book on this subject by Chopra and Marfurt, edited by Hill, later in the year. Both the book and the DISC present a variety of methods and numerous examples of their application. These techniques are highly mathematical. This TLE *paper is devoid of both examples and equations and presents the methods in a compact fashion. It is designed to give DISC attendees, prior to the course, a better grounding in the methods and to facilitate comparing those methods with each other.*

from that similarity that our eye so easily notes along that horizon.

If we observe a depth slice of the data, such as Figure 3, we lose the sense of the lateral similarity. That depth slice does not represent enough data. For this reason, we will consider a slab of the rock, not just an interface, for our similarity analysis. What is a slab? As an example, Figure 7 shows a slab of petrified wood. In our seismic data, the slab will have a thickness of 5–40 ms.

We wish to determine the lateral similarity in a slab of seismic data. But, just what do we mean by similarity?

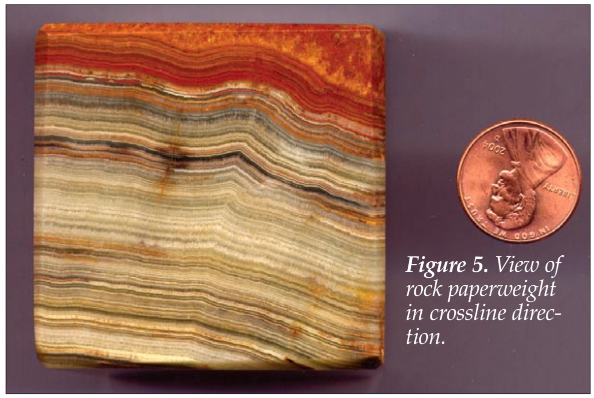

Figure 5. View of rock paperweight in crossline direction.

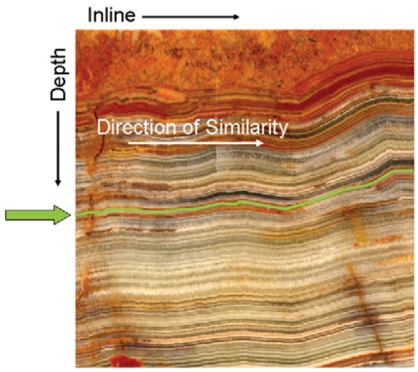

Figure 6. View of rock paperweight in inline direction. Same view as in Figure 4. Arrow points to green line which is a picked horizon.

A classification example. Figure 8 is a photograph of a collection of household items. Based on their shapes, which items are most similar to each other? Which items are most dissimilar? This is a pretty trivial assignment. Your mind instantly notes the similarity of these items not only to each other, but also to standard objects that are in your mind. That is the reason that you can identify object C as a nut, even though you have never seen this particular nut before. You used your mind's standard nut as your classification (naming) tool because the photograph of this nut had greater similarity to your standard nut than any other item that resides in your mind's naming catalog.

The four bolts in Figure 8 (A, D, F and G) are sufficiently similar to each other to all fall in the "bolt" category. Items D and F are the most similar, even though they cannot be absolutely identical to each other. (After all, there will be slight variations in the manufacturing process.) When we see these items, we also think of their applications or, in other words, how they would be used. If asked, based on our experience and knowledge, we could also provide an explanation about their manufacturing history, starting with the extraction of the iron ore. Our visual classification assists us in placing together items related to similar processes. That is also the goal of the classification of the individual traces in seismic data. We will hope that traces that we classify in similar categories will have been created through a similar manufacturing (or for geoscientists, deposition, deformation, or diagenetic) process.

Figure 7. A slab of petrified wood.

What have we learned from this foolish exercise with these household objects? We have learned that our mind automatically classifies these household items against each other and also against external criteria that we keep in our mind. Even without knowing the names of these objects, we can determine which are most similar to each other, and in what regards. For example, we may classify them as being similar in length, similar in color, rusty or shiny, or similar in apparent age. These measures are "attributes" of these simple household items. We also can apply an external standard, such as a ruler, to estimate the actual lengths. We can apply an external standard to conclude that we have a nail, a few nuts, and a loose screw. The point is that we can judge similarity in a relative fashion, one object to the next, or we can judge similarity against an external standard.

Crosscorrelation as a similarity measure. Now we turn to seismic data. We wish to classify windows of our seismic traces based on their similarity. (Later sections will answer the question "Similarity to what?") Because of the overwhelming number of traces in a 3D survey, we require a computer-based, numerical method. Since our seismic traces are time series, mathematicians inform us that the appropriate method is crosscorrelation which takes advantage of the power of addition and multiplication.

We can use the mathematical recipe of crosscorrelation to determine how much trace 1 looks like trace 2 in Figure 9. In this example, trace 1 is identical to trace 2. The crosscorrelation recipe states that we should first pairwise multiply each of the amplitude time samples of trace 1 with trace 2, producing a new, intermediate trace labeled "product." Because traces 1 and 2 are identical, the product trace contains the square of the amplitude values in either trace 1 or trace 2. All values in the product trace will be positive or zero. (Note that the multiplication of the negative values in trace 1 by the negative values in trace 2 produces a positive product.) The next step is to sum all values in the product trace. This produces a single number that is proportional to the crosscorrelation. Because all of the values in the product trace are positive, the sum is a large, positive

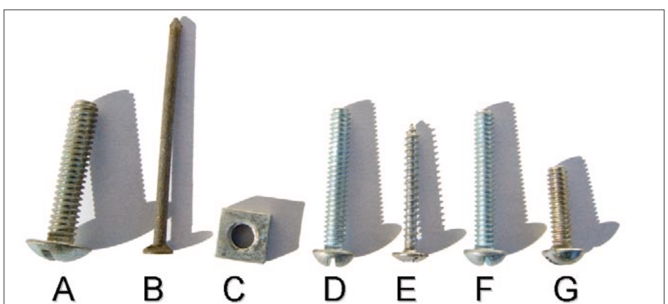

Figure 8. Selected collection of household objects.

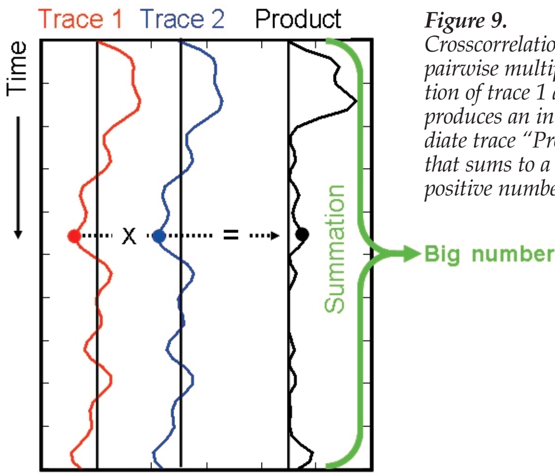

Figure 9. Crosscorrelation: pairwise multiplication of trace 1 and 2 produces an intermediate trace "Product" that sums to a big positive number.

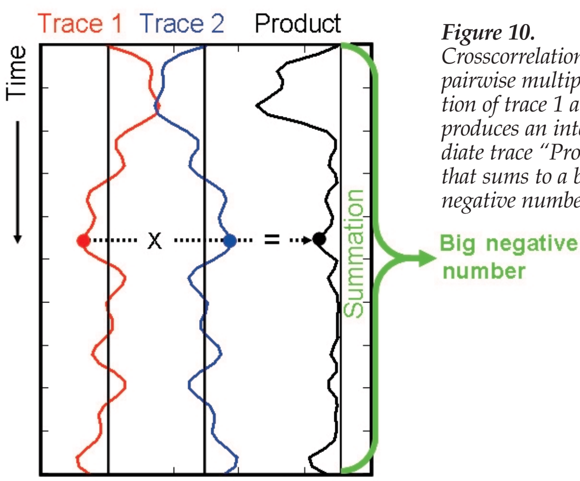

Figure 10. Crosscorrelation: this pairwise multiplication of trace 1 and 2 produces an intermediate trace "Product" that sums to a big negative number.

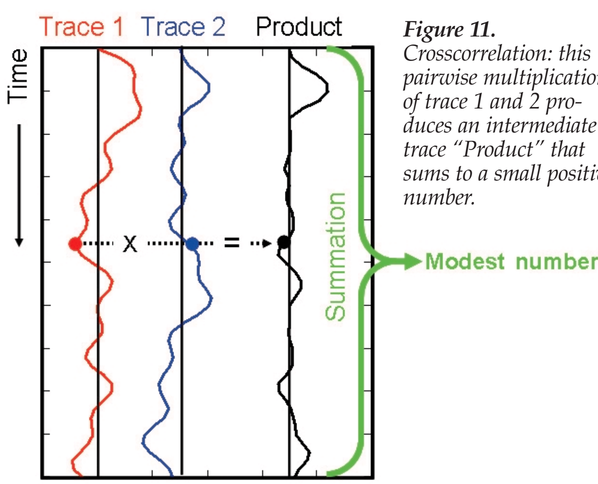

Figure 11. Crosscorrelation: this pairwise multiplication of trace 1 and 2 produces an intermediate trace "Product" that sums to a small positive number.

number.

We now investigate the situation for which trace 2 is as different from trace 1 as possible; i.e., the data in trace 2 are the negatives of the data trace 1 (Figure 10). In this situation, the product trace contains only negative values because each pair-wise multiplication produces a negative value or zero. So, the summation of the product trace produces a large, negative number. This negative number is the crosscorrelation coefficient between traces 1 and 2.

We now consider the middle road. In Figure 11, trace 2 is randomly different from trace 1 and pairwise multiplication produces both positive and negative values in the product trace. In this particular case, the summation of all values in the product trace produces a small, positive value for the crosscorrelation coefficient. With a randomly different trace 2, the crosscorrelation coefficient could have been a small negative number.

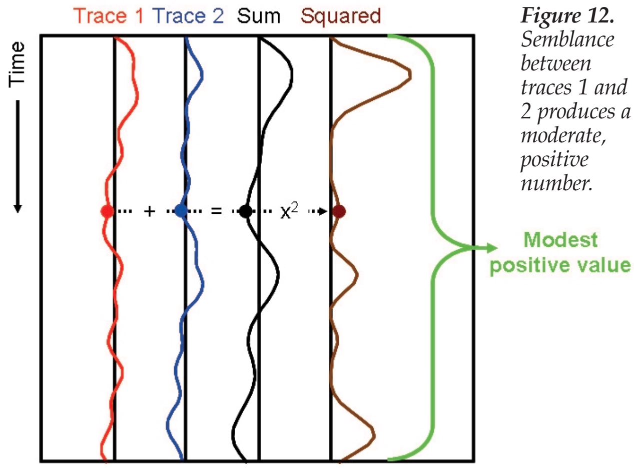

Figure 12. Semblance between traces 1 and 2 produces a moderate, positive number.

While it may not be apparent, crosscorrelation is not as smart as our eye. If we have two identical traces, but one is time-shifted with respect to the other, then the crosscorrelation coefficient will be smaller than if they had not been time-shifted. Crosscorrelation is a bit dumb, in that regard. While our eyes would have no problem discerning that these traces were time-shifted, identical traces, crosscorrelation is fooled. This is the reason that it is advisable to remove the local dip from our data in advance of crosscorrelation.

In order to keep things tidy, mathematicians sometimes add an additional step to this crosscorrelation recipe. After summing the product time series, they normalize that summation by dividing by individual summations of the two input traces. This means that the crosscorrelation coefficient is one for identical traces, minus one for traces that are the negative of each other, and small values, close to zero, for traces that are randomly dissimilar. These are termed "normalized" crosscorrelation values and the extra step in the recipe is termed a "normalization" step. The normalized crosscorrelation will be insensitive to changes in the scaling of the amplitudes of either input traces, as long as the amplitude scaling factor is a positive number.

Semblance as a similarity measure. Another measure of trace-to-trace similarity is semblance (Figure 12). Take a minute to compare the operation in Figure 12 with the operation in Figure 11. Both compare trace 1 to trace 2 and both have a step that produces a new, intermediate trace. In crosscorrelation, trace 1 and trace 2 were multiplied together. In semblance, trace 1 and trace 2 are simply added to each other. In addition to replacing a multiplication with a summation, there is an additional difference. In the semblance calculation, the intermediate trace is squared before the summation step. As was the case with crosscorrelation, if the two traces are identical, then their summation is large. However, if the two traces are of opposite sign, then their summation produces an intermediate trace with all zeros. Unlike crosscorrelation, the squaring operation in the semblance calculation ensures that the semblance is never negative.

There are two additional items to consider in understanding the semblance calculation. First, the semblance output is generally normalized by dividing by the summation of the squares of the individual traces. In spite of this, mathematicians inform us that the semblance remains sensitive to changes in the overall amplitude scale factor of the traces. Second, this semblance definition is easily generalized for the case of additional input traces. For 3D data with a central trace and eight adjacent, surrounding traces, the first step in the semblance calculation sums (stacks) all nine traces.

So, where are we? We now have a couple of methods that we can use in a computer to create numerical scores which measure the degree of similarity between two traces. If we have many traces, the semblance calculation simply adds them together as the first step in the calculation. By contrast, because crosscorrelation is defined for only two traces, the inclusion of more traces implies that the different crosscorrelation values obtained between pairs of traces must be combined.

Now we are getting somewhere. We have a couple of tools, crosscorrelation and semblance, to quantify trace similarity and are now halfway to the end of the story. With a crosscorrelation coefficient for every data trace, we have replaced our finite-thickness slab of seismic wiggle traces with a horizon of crosscorrelation coefficients displayed at the X, Y locations of the original traces. The crosscorrelation can be between seismic traces that reside in the data or between the traces and a selected external standard. Just as in classifying the objects in Figure 8, we have the freedom to select internal similarity comparisons or comparisons to an external standard. The next sections review both approaches.

Internal similarity. We will first consider crosscorrelations between traces that reside within the seismic data. This case has two subdivisions. The first determines the similarity to nearby traces. The second determines similarity to more distant traces.

Similarity to immediate neighbor traces. Figure 13 shows the lateral location of three traces, A, B, and C. We obtain the normalized crosscorrelation coefficients between traces A and B and also between traces A and C. These normalized crosscorrelation coefficients are then combined (the square root of the product) to provide a representative crosscorrelation between trace A and a pair of its neighbors. The inventors termed this measure of similarity *coherence*. In order to remove the effect of local dip, the crosscorrelations are performed with a variety of relative time shifts between A and C and A and B and the maximum value of the respective crosscorrelations is used in the recipe. With this modification the procedure can track changes in the dip. (This modification is applicable if you had forgotten to flatten to your horizon of interest before the coherence determination.) Because this method uses a normalized crosscorrelation, the crosscorrelation can range only between -1 and +1, and is independent of the actual amplitude scaling of the trace. (The squaring operation in the recipe converts negative crosscorrelation coefficients to positive numbers.) Thus, this coherence method is blind to amplitude scaling differences, trace-to-trace. Of course, in comparison to trace A, if one wiggle in trace B in the analysis slab interval has a change in amplitude and another wiggle on the same trace does not, then traces A and B have different shapes and the crosscorrelation coefficient will also change.

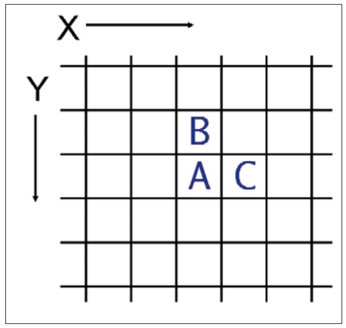

Figure 13. *View of X, Y locations of three seismic traces.*

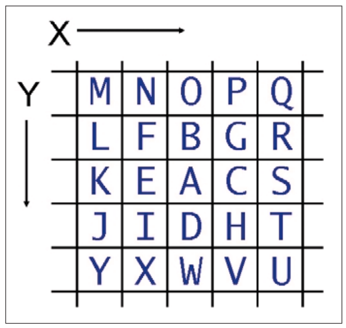

Figure 14. *View of X, Y locations of 25 seismic traces.*

Dip determination. For this purpose, the crosscorrelation calculates the similarity along different dips in the data and selects the dip that has the greatest similarity. From the extracted dips, we can then calculate the reflector curvature. For 3D data, the algorithm supplies curvature in two perpendicular directions to reveal a local bowl, dome, saddle, plane, etc.

Semblance-based similarity. You may well ask, when viewing Figure 13, "Why not include more traces in the calculation?" And, yes, that can be done. Figure 14 illustrates just such an increase in the number of traces. It would be possible to obtain the crosscorrelations between traces A and all other illustrated traces, B through Y. There are innumerable ways to combine these crosscorrelation values. The values could be averaged, or a distance-weighted average could be used. Or, the algorithm could take the square root of the sum of the squares. However, this is not common practice. Instead, it is customary to use semblance, and not crosscorrelation, to reveal the degree of trace-to-trace similarity. If signal-to-noise is a problem, then increasing the number of input traces might be the solution.

In contrast with the crosscorrelation-based similarity measure, the semblance-based measure is sensitive to relative changes of the scale of the trace amplitudes in addition to differences in the shapes of the collection of traces.

Dissimilarity measure. A variation of the semblance-based algorithm determines how dissimilar the traces are by subtracting each trace from the average of all traces. This recipe produces what the mathematicians call the variance. Appearances to the contrary, the variance does not provide any new information because mathematicians inform us that the variance is simply equal to one minus the semblance.

Eigenstructure-based similarity. This is a doozy of a recipe. All possible crosscorrelations between all traces provide the only input to the method. To be explicit, following the labeling in Figure 14, we first crosscorrelate trace M with itself, then M with L, then M with K, and so on. Then we crosscorrelate N with itself, N with F, etc. We end up with every possible two-trace crosscorrelation. Proceeding through a series of steps that almost belie belief, we obtain a single number that is a measure of the similarity of all traces with each other. So, what's the good news about this? There is theoretical evidence that the output number is more robust against the presence of random noise in the input traces in comparison to the multitrace, semblance-based measure. In addition, the eigenstructure-based similarity value is independent of the application of positive or negative amplitude scale factors.

Unfortunately, we cannot think of a simple cartoon that illustrates determination of the eigenvalue similarity from the collection of the input crosscorrelations. However, the next three illustrations might atone for that deficiency. Consider the simplest case, two traces. The left side of Figure 15 shows a red trace and a blue trace. Visual inspection indicates that the traces have a high degree of similarity. The large value of their crosscorrelation (0.99) is noted. The right side of the illustration shows a crossplot between the amplitudes of the red and the blue traces. If the traces were identical, then the amplitude values would be a straight line. The traces are not exactly identical, so there is a slight scatter of amplitude values from a straight line. Also shown in this graph is a coordinate system rotated for maximum alignment of one of the two perpendicular axes with the trend of points. The calculation of an "eigenvector" from the crosscorrelations provides just this rotation. The associated, largest value of the eigenvalue is a measure of the scatter of the points to this rotated coordinate system. This eigenvalue, divided by the sum of all output eigenvalues, gives the eigenvalue similarity. The eigenvalue similarity in this case is also large, 0.99.

Figure 16 has the same format as in the previous figure but with a decreased similarity between the red and the blue traces which produced increased scatter perpendicular to the "longer" rotated coordinate. Notice that the eigenvalue similarity has decreased.

Figure 17 shows an even more extreme case and an even smaller value of the eigenvalue similarity.

With an increased number of traces, the story is the same, simply more difficult to visualize. With nine traces, for example, we would have a nine-dimensional coordinate system replacing the two-dimensional coordinate systems shown on the preceding series of figures. Imagine a nine-dimensional coordinate system filled with points representing the amplitude values at a given time along each of those nine traces. For example, if the points are scattered about in that nine-dimensional super-volume, then their eigenvalue-based similarity will be low. If they are along a straight line in this nine-dimensional space, then the eigenvalue-based similarity will be large. With the additional traces, we would have the same story as the illustrated two-trace case, just a bit more difficult to visualize.

Nearby dissimilarity. Instead of highlighting trace-to-trace similarity, we could have started with just the opposite goal, that of highlighting trace-to-trace dissimilarity. Look again at Figure 12, the figure that illustrates a recipe for emphasizing trace-to-trace similarity through semblance. Now consider Figure 18. It is almost identical to Figure 12 with one significant difference in the first step. In this case, one trace is subtracted from the other. Then, their difference is accentuated by squaring. The output is the sum of that squared trace. Mathematicians term this a *Sobel filter*. The output of this operation will depend upon scaling differences in the adjacent traces, in addition to waveform differences.

Distant similarity. We have now considered many of the similarity measures (and a dissimilarity measure) between and among a nearby neighborhood of traces. Now we turn to the similarity to other traces from the volume.

Known calibration. Imagine that we already have well control in the seismic volume and have found a horizon that contains our desired pay. We know the horizon of this discovery and the location of unsuccessful, dry wells in the same horizon. We notice that the waveform at the discovery horizon appears somewhat different from the waveform at our dry hole. Of course, we want to find more discoveries. So, we can use crosscorrelation as our tool to determine the similarity of all traces to the trace at the discovery well. From that, we produce a correlation, i.e. "drill-here," map.

Unknown calibration. By contrast, assume that we do not yet have a discovery well; we have not yet unlocked the key to this region. We do observe that the waveform does change laterally along our prospective horizon. We could use a simple amplitude map at that horizon, but it does not fully represent the lateral changes in the time-dependent waveform. In the absence of a well, all is not lost, however. We can use the data themselves to provide us with a collection of representative waveforms (let's say ten waveforms) at the prospective horizon and then crosscorrelate each waveform against our extracted slab of seismic data. Having arbitrarily numbered each reference waveform (1-10), we can create a map that shows, as a function of X, Y which particular numbered waveform is most like the numbered, reference waveform at all X, Y locations. How do we determine this best-match waveform? We will use crosscorrelation, for

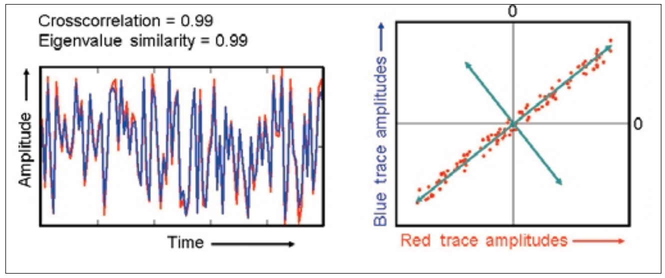

Figure 15. *Two almost identical traces and their amplitude crossplots.*

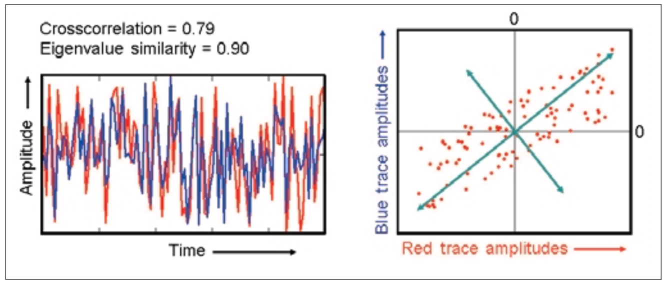

Figure 16. *Two somewhat identical traces and their amplitude crossplots.*

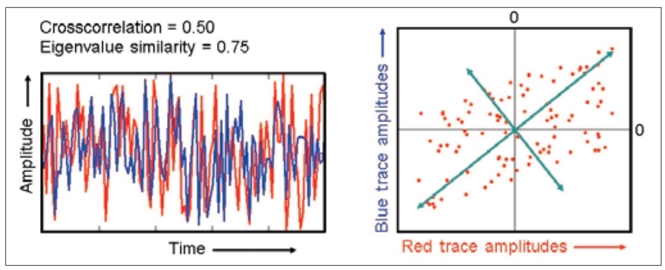

Figure 17. *Two dissimilar traces and their amplitude crossplots.*

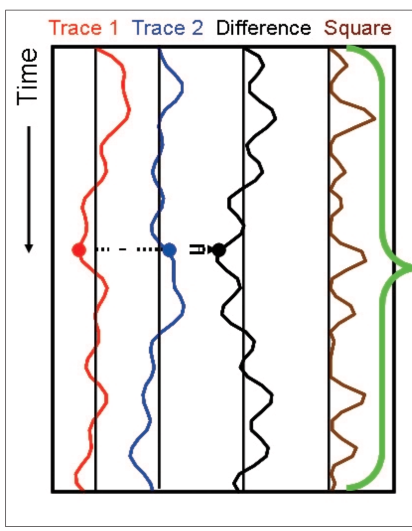

Figure 18. *Sobel filter. The sum of the trace difference squared.*

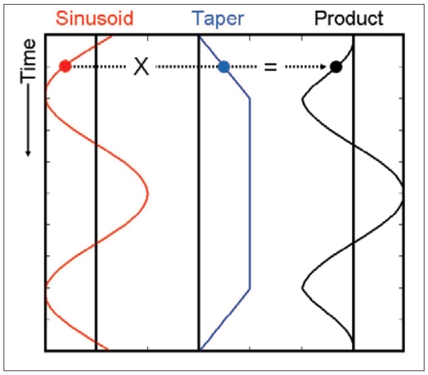

Figure 19. *Tapering a sinusoidal function.*

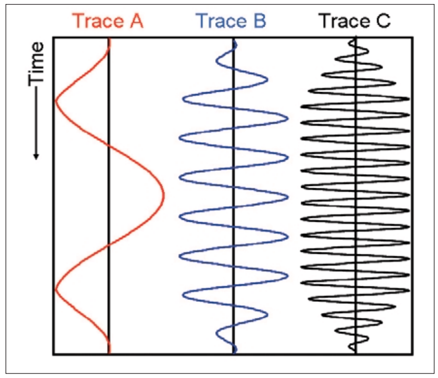

Figure 20. *Tapered wavelets at three different frequencies.*

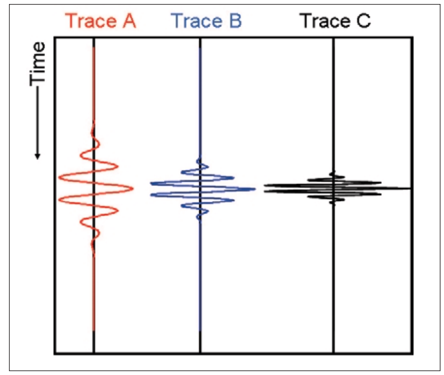

Figure 21. *Wavelet transform wavelets.*

example, to score the competing reference waveforms against each other. When we are done, instead of producing an X, Y map of the winning numbers, we can represent each of the waveform numbers by a distinct color. This will provide a visual display of the similarity of the traces in the selected slab to this set of data-extracted, idealized traces. If we are lucky, we may observe that this colored map might reveal geologic shapes among the different colors.

External similarity. The previous section explained the use of the data itself to create its own set of waveform standard. By contrast, we could use an external standard, a "ruler," in our trace similarity calculation. Thinking back to the collection of household items in Figure 8, remember that you instantly named each item. Assigning the names of nail, screw, and bolt was an illustration of the application of an external similarity test.

At present, there are two popular sets of external, standard waveforms used as "rulers." These standard waveforms are very similar because both are sinusoidal in nature. The use of these sinusoidal waveforms is termed *spectral decomposition*.

Tapered, windowed wavelets. Because the traces in our data slab are oscillatory, a set of sinusoidal functions provides an excellent candidate for our set of standard traces. We create a series of sinusoidal functions at different frequencies, and then determine the crosscorrelation coefficients between these sinusoidal functions and the traces in the slab of seismic data. The crosscorrelation coefficient between each standard trace and the seismic trace is in proportion to the amount of that selected frequency that resides in each trace. (This is something that M. Fourier determined quite a few years ago.)

Why would we want to crosscorrelate against a series of cosine waves? If the series of reflectors within our slab move closer to each other, the reflectors start to look more like a high-frequency sinusoid instead of a lower-frequency sinusoid. Thus, the crosscorrelation coefficient will increase for the high-frequency sinusoids and, relatively speaking, decrease for the low-frequency sinusoids. From this we might be able to see subtle thickness changes such as in Figure 4.

There is, however, one fly in our ointment. Our slab is not infinitely thick. It is a finite-thickness slab. In fact, to improve vertical resolution of our analysis, we use a thin slab so we can determine the characteristics of a small number of interfaces or a thin formation. The short-length slab traces will produce almost identical crosscorrelation values when correlated with the collection of sinusoidal-frequency standard traces of neighboring frequencies. To consider a concrete case, pretend a trace from our seismic slab is a 30-Hz sinusoidal function. Because the slab trace is truncated in time, its crosscorrelations with the 27-, 28-, 29-, 30-, 31-, 32-, and 33-Hz sinusoidal-frequency standard traces produce almost equivalent crosscorrelation coefficients. In fact, crosscorrelations with even low- and higher-frequency standard traces reveals a smoothly varying, frequency-dependent change in the values of the crosscorrelation coefficient. The thinner the slab in time, the greater fuzz in our determination of the frequencies that reside in that slab. (Yes, we geophysicists have our own uncertainty principle.) Mathematicians tell us that we can minimize that fuzziness if we taper the sinusoidal standard traces. Figure 19 cartoons the tapering. The trace on the left is a single frequency sinusoidal function. The middle trace is the tapering operator. The zero–to-zero width of the taper is the time thickness of our slab. (Note, we could, equivalently, taper the value of the amplitudes toward the bottom and top of the slab of seismic traces.)

The right trace is the resultant tapered wavelet. Figure 20 is a cartoon of three different windowed/tapered sinusoidal functions that we will use as our external standards. It is these windowed/tapered traces that are crosscorrelated against the traces in the slab. As a result of these crosscorrelations, we obtain one set of crosscorrelation coefficients for each sinusoidal frequency. As before, we can display these crosscorrelation values in an X, Y sheet for our selected frequency. Because of the large number of sinusoidal frequencies we can select for the external standard, it is possible to produce more crosscorrelation sheets than input time samples from the original slab. Of course, this extra information comes at a price. The output crosscorrelation sheets, one frequency to the next, will be quite similar and, hence, do not supply us with independent, new information.

As a review of your comprehension of these concepts, we invite you to think through the following conclusion: the unnormalized crosscorrelation of the slab traces with a very low-frequency standard trace is, within a scale factor, the numerical average of the amplitudes within the seismic slab.

Wavelet-transform wavelets. Knowing that it is desirable to use tapered wavelets, such as in Figure 20, other practitioners choose to use the set of external-standard wavelets that come from what mathematicians term the *wavelet transform* method. Wavelet-transform wavelets are manufactured in a fashion similar to the tapered sinusoids (Figure 19), with one critical difference—the tapered window does not have a constant length but, instead, the length depends upon the sinusoidal frequency of Figure 19's first trace. Through the appropriate mathematics for the tapering function, each series of wavelet-transform wavelets appears to be a time-stretched version of each other. Figure 21 shows three such wavelet-transform wavelets.

Table 1. *Uses of seismic attribute methods*

Slab attribute	Measures	Examples
Dip and azimuth	Lateral changes in reflection time within a specified window	Creation of dip map if individual horizons difficult to pick Revealing karst-enhanced collapse features Highlight lineaments, folds, and well-focused faulting Revealing out-of-plane dips in vertical slice of data Classify reflector convergence and divergence in 3D volume Forms the input for curvature and reflector shape computations Forms reference for coherence, reflector amplitude gradients, and structure-oriented filtering
Coherence (semblance and variance)	Trace-to-trace similarity (including changes in trace amplitudes)	Reveal thick channels that create lateral changes in waveform Reveal faults in time-slice parallel or perpendicular to stratigraphic bedding Decrease the sensitivity to differences in wavelet phase and frequency content when comparing multiple 3D surveys Slumps and scours that produce lateral discontinuity in average reflection coefficients (Pinnacle) reef boundaries
Coherence (crosscorrelation and Eigenstructure)	Trace-to-trace similarity (insensitive to changes in trace amplitudes)	As above Insensitive to lateral changes in amplitude
Curvature	Changes in dip	Curvature proportional to fracturing? Faults with throw less than 1/4 wavelength not seen on coherence Faults separating rotated blocks Channels if associated with differential compaction (not seen in Tertiary Gulf of Mexico) Unbroken flexures Ability to filter for long-wavelength changes in curvature to map poorly focused faults Track edge of reef
Shape index	Curvature in two perpendicular directions	Collapse features, domes, ridges, valleys
Amplitude gradients (not mentioned in text)	Lateral differences in amplitude	Faults and fractures Lateral changes in thin channel tuning, revealing channels where no lateral change in waveform, only amplitude Slump features Scours that produce lateral discontinuity in average reflection coefficients Channels (sometimes thinner than seen in coherence) Edges of reef thinning
Spectral decomposition	Similarity to sinusoidal functions	Changes in thickness of channel by frequency-dependent tuning Edges of valley fill Stages in valley fill noted by changes in waveform amplitude (but, perhaps not, changes in waveform shape) Improving gas bright-spot standout (at high frequencies) Reveal frequency-dependent absorption due to thick shallow gas reservoir Faulting Thin reef reservoir thickness
Continuous wavelet transform spectral decomposition	Similarity to finite-length "wavelets" (wavelets are mathematically defined)	Improved visibility of channels Observe lateral changes in reservoir thickness Reveal frequency-dependent absorption due to thick shallow gas reservoir

Summary and conclusion. The transition from 2D to 3D seismic data brings the challenge of 3D visualization. Because we are investigating layered media, our geophysical community has been able to respond with a series of trace-to-trace, horizon-tracking, similarity measures. These similarity measures quantify the degree to which a given trace looks like (or differs from) a second trace. For the most part, the similarity measure is a crosscorrelation coefficient, or a near relative of that recipe. For an internally derived reference, the second trace in the crosscorrelation can be a near-neighbor trace or a standard trace created from the data itself. For an externally derived reference, the second trace can be created from sinusoidal functions of a series of neighboring frequencies.

Through the use of these similarity measures, we convert a 3D slab into a 2D sheet that hopefully contains relevant information from the 3D slab.

This review highlights the common elements of the different slab-based similarity analysis tools. Many algorithms condense the seismic slab to a sheet of numbers, formed by a weighted average of the trace amplitudes. In some cases, an internal standard, such as a neighboring trace, determines the weighting function. In other cases, sinusoidal functions serve as the weighting function. This similarity of many of the processes accounts for the approximate visual similarity in the output of the application of some of these tools.

So, finally, we come to the bottom line: How do interpreters actually use these tools? Table 1, derived from examples in the upcoming book by Chopra and Marfurt, lists the most common (at the moment) applications. TLE

Corresponding author: stevejhill@earthlink.net