Detection of seismic discontinuities

ABSTRACT

In seismic images, the opaque oscillatory nature of the ever present sedimentary layers often hide the boundaries between sedimentary packages, such as faults or river channels. Interpreters painstakingly identify these boundaries by inspecting individual planes of the image volume. Discontinuity attribute maps computed from image volumes depict these boundaries and greatly simplify and accelerate seismic image interpretation. Standard edge enhancement fails to compute useful discontinuity maps, since the layers within a sedimentary package show similar rapid amplitude changes as the sought after package boundaries. Instead, the correlation of a local image region – *patch* – and its best-fitting plane wave delineates discontinuities that are characterized by a very rapid, pixel-to-pixel change in amplitude. The correlation attribute fails to delineate comparatively smooth discontinuities such as fault zones or fault reflections. Finally, locally estimated prediction-error filters fail to distinguish discontinuities and other image features. A prediction-error filter predicts and removes nearly all features in a migrated seismic image volume. Both, the misfit of a local plane wave and the prediction error only indicate if a given local image region contains a discontinuity or not. Consequently, the resolution of both approaches is limited to the size of the local image patch, which itself is somewhat arbitrary but subject to limiting tradeoffs. Future development ought to lead to a more fundamental classification of seismic image features. Discontinuity maps are only a special case of a potential automatic interpretation system.

In the following chapters, I compare several methods to compute discontinuity attribute maps from seismic image volumes. These discontinuity maps aim to delineate boundaries of sedimentary packages, such as fault planes, river channels, and salt flanks. Discontinuities maps accelerate and improve the structural interpretation of traditional seismic image volumes. Figure 1 shows a time slice of a Gulf salt dome image and a corresponding discontinuity image – incidently my best result. In this chapter, I introduce the detection of seismic discontinuities by mentioning related publications, discussing the causes of seismic image discontinuities, outlining the investigations of the subsequent chapters, and presenting the three test cases that I use in this study.

Literature

Traditionally, seismic processing back-propagates a surface recorded wavefield to its subsurface reflectors and thereby creates a *wiggly* image of the subsurface. The back-propagated wavefield enables an interpreter to derive a model of the structural and stratigraphic geology of the imaged subsurface (**???**). The interpreter integrates the initial structural image with additional geological and geophysical information to a complete geological subsurface

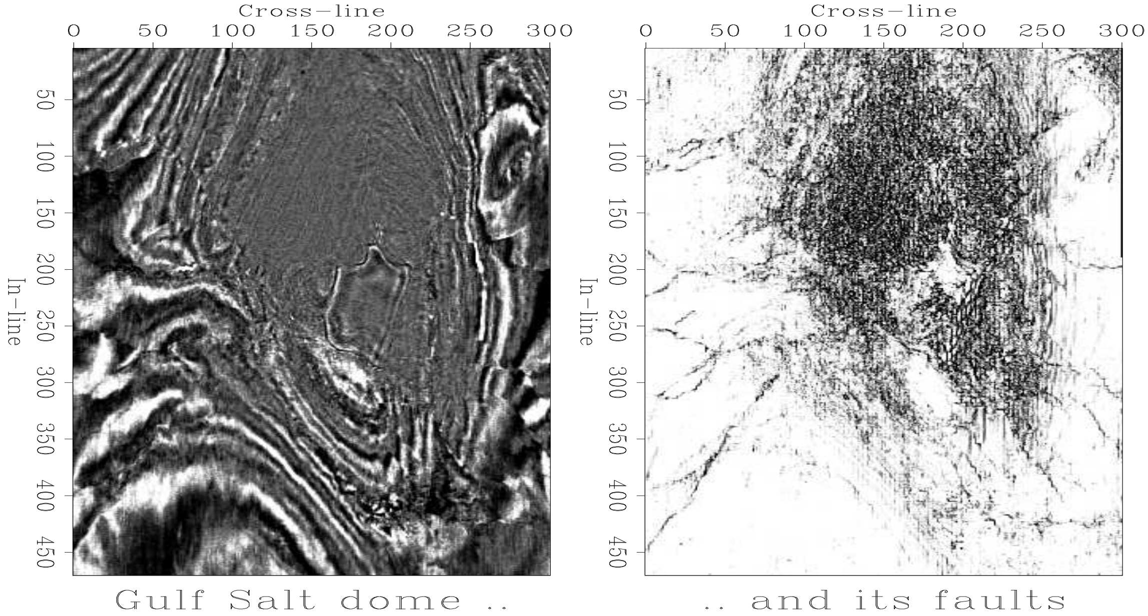

Figure 1: Gulf salt dome and its discontinuities. The discontinuity image at the right delineates the faults of the original seismic image at the left. title [CR]

model.

In the last decade, geophysical researchers have attempted to extract information beyond the wiggly, structural image from high-quality seismic recordings. Wavefield inversion (Bleistein, 1984; Tarantola, 1987; Weglein, 1989) estimates physical rock properties, such as local P-, S-wave velocity, and density. Amplitude-versus-Offset analysis relates the prestack seismic data to pore fluid contents or hydrocarbon indicators (**?**). Seismic attribute mapping empirically correlates pre- or poststack seismic data characteristics, e.g., instantaneous frequency, to rock properties (Taner et al., 1979; Sonneland et al., 1990). Many studies combine the various approaches – inversion, AVO, seismic attributes – with other data – e.g., well log data – to extract improved lithological and petrophysical subsurface information (**?**). The overall goal of this research is to generate continuous, reliable, geological models of the subsurface.

While much work has concentrated on widening the application of seismic data to extract additional information, limited research has been done to improve the extraction of the traditional structural interpretation. Today's interpreters explore huge three-dimensional image volumes by displaying them on graphics workstations (**?**Sonneland, 1983) and by plotting traditional paper sections. These media force interpreters to identify three-dimensional features – river channels, faults – on two-dimensional cross-sections of the image volume. Because of the wavefield character of the seismic image, the interpreter lacks a depth view of the features of interest. To improve picking accuracy, interpreters may use seismic event picking software that evaluates a picked image location by local volume correlation.

Amoco's fault detection technique, *Coherency Cube* (Bahorich and Farmer, 1995), aims to help interpreters by transforming a traditional seismic image volume to a discontinuity image volume. Marfurt et al. (**?**) published an improved scheme. In principle, both scheme's resemble the plane-wave misfit, I introduce in this paper. However, I replaced Marfurt's compute-intensive search for the plane-wave dip by a simple analytic estimation. Coherency Cube generated much interest and, subsequently, other researchers (**?**Gersztenkorn and Marfurt, 1996) presented similar techniques. Bednar (**?**) presented a discontinuity attribute based on least-squares dip estimation. Today many commercial seismic processing and interpretation packages include a discontinuity attribute process; their methodologies are usually confidential or patented.

Fundamentally, image enhancement amplifies the pixel amplitudes of the target features relative to the amplitude of the undesired background features. Transform methods, such as Fourier, Wavelet, Hough, or Principal Component decomposition, explicitly separate the image into components that can then be recombined using target enhancing weights. Prediction methods, on the other hand, assume the absence of the target feature and produce high amplitude output for image pixels that deviate from that assumption. However, transform and prediction methods are two modes of thinking rather than two distinct approaches: the Laplace operator, a standard edge enhancement technique (Jain, 1989; **?**), can be thought of as a prediction of the central pixel amplitude derived from its surrounding pixel values. Alternatively, the Laplace operator is often implemented by Fourier transform, in which case it is best thought of as amplifying the high-frequency image components.

Discontinuity, picking, interpolation, and data quality

Discontinuity attribute, seismic reflector picking, interpolation of seismic data, and the assessment of image quality all require a mathematical model of seismic data. A successful model for one application might work for other applications. For example, I derived the discontinuity attribute that is based on the best-fitting plane wave from a dip-picking scheme by Claerbout (1992). Before using prediction error to detect faults, I (1995) used it to interpret seismic data. Claerbout applies a similar prediction error approach to detect lapses in data quality.

Over the years, I painfully learned, however, that the characterization of seismic data greatly varies among its domains, prestack and image domain as well as two-dimensional and three-dimensional domain. A method that works in one domain does not necessarily succeeds in another. For example, prediction-error filters interpolate the superposed plane waves of prestack data, but fail to remove the adjacent plane waves separated by a fault. Wavelet compression succeeds in the high-dimensional, redundant prestack domain but fails for two-dimensional seismic images.

Future

Interpretation of seismic subsurface images can and ought to be aided by computer processing. The delineation of faults is a step towards general automatic information extraction from seismic images. Automatic seismic image analysis would ease the tedious and time-consuming extraction of geological information from seismic image volumes by human experts. Additionally, automatic analysis could objectively and reproducibly compare the effect of competing seismic image processes on the interpretation of seismic data. However, artificial intelligence experiences great difficulties in the development of visual scene recognition and interpretation (**?**).

For continued research into automatic seismic image analysis, I advocate (1) a classification framework for seismic features and (2) a sophisticated segmentation scheme. I wonder if basic statistical measures – mean, mode, variance, autocorrelation, and grey level histograms – can classify seismic image features and their boundaries – layered sediments, salt, base rock, off-lap surfaces, on-lap surfaces. If such a classification is feasible, the subsurface could be separated into amorphous segments of similar, stationary statistics and, consequently, regions of identical geology. Interpreters of remote sensing data successfully use similar classification and segmentation algorithms (**?**).

My investigations

In this chapter I explore three discontinuity attributes: standard edge enhancement, plane-wave misfit, and prediction error.

I illustrate each approach on one synthetic (Figure 4) and two seismic images (Figures 7 and 11). The first seismic image depicts a Gulf salt dome. The faults of the image are sharp and well defined. Bahorich et al illustrates the original coherency attribute at the same image. The second seismic image shows a North Sea horst and is the more challenging test case. The faults are smooth and difficult to find in the seismic time slice. Bednar used the same image in the publication that describes his discontinuity attribute.

The Table 1 lists the various methods that I include in this publication. The standard edge enhancement techniques fail to differentiate sedimentary beds and boundaries of sedimentary packages (Figure 16). However, two-dimensional edge enhancement within a time slice takes advantage of the dominant horizontal dip of sedimentary strata and yields reasonable fault maps (Figure 23).

Of all my discontinuity attributes, the trace correlation of an image region and its locally best-fitting plane wave consistently yields the clearest fault maps (Figures 33 and 34). Other measures of plane-wave misfit that I tested generated only inferior fault images (e.g., Figure 27).

The algorithm of my plane-wave correlation is, similar to the computation of the commercial coherency attribute by Bahorich and Farmer (1995). Both processes certainly generate a similar discontinuity map of the Gulf salt dome. My plane-wave estimation is related to Symes' (**?**) differential semblance. I derived the plane-wave estimation by generalizing a dip estimation scheme by Claerbout (1994). Bednar derives his discontinuity attribute from Claerbout's dip estimation, too.

Prediction-error is an unreliable and noisy discontinuity attribute. The discontinuity map of the salt dome (Figure 44) shows the radiating faults, and even resolves features within the salt. Unfortunately, the image is obscured by noise. Even worse, the prediction

Label	Name	Computation	Output format
Gradient methods			
Grad	Gradient magnitude	$\|\nabla g(x)\|$	scalar at each pixel
Lap	Laplace	$\nabla^2 g(x)$	scalar at each pixel
Dip misfit methods			
DXG	Cross-product	$\mathbf{r} = \mathbf{p} \times \nabla g$	3 component vector at each pixel
DXGLN	Magnitude of r	$\|\mathbf{r}\|$	scalar at each pixel
DXGAA	Backprojection of r	$(\mathbf{p} \times \nabla) \cdot (\mathbf{p} \times \nabla) g$	scalar at each pixel
Waveform misfit methods			
PPC	Correlation of estimate and original	$Cor_z[g(\mathbf{x}), f(\mathbf{p} \cdot \mathbf{x})]$	scalar for each trace of patch
Prediction-error methods			
Pef3d	3-D PEF	$A\, g$	scalar at each pixel
XPef	3 2-D PEF	$\epsilon = [A_x, A_y, A_z] g$	3 component vector at each pixel
XPefLN	Magnitude of ϵ	$\|\epsilon\|$	scalar at each pixel
XPefAA	Backprojection of ϵ	$[A_x, A_y, A_z] \cdot \epsilon$	scalar at each pixel

Table 1: Discontinuity attributes. The table lists the various discontinuity attributes that I explore in this thesis. An attribute's label identifies the corresponding names of Figures. The norm $|\cdot|$ is the standard vector norm. The vector $\mathbf{p}$ is an estimate of the local plane-wave normal. The operator A represents a 3-D prediction-error operator. The operators A_i are two-dimensional versions.

error filter predicts and removes almost all features of the smoother North Sea horst image (Figure 45).

Prediction-error images tend to be white (Jain, 1989), since the filter removes the predictable correlated components of its input. Claerbout (1992) introduced a combination of two two-dimensional filters that detect and suppress a three-dimensional plane-wave. I added a third prediction-error filter, which generalizes Claerbout's intuitive, original formulation. The three prediction-errors are linked to the cross-product expression of my plane-wave formulation.

The prediction-error filters' removal of all image features in the North Sea case suggests limiting the predictive power of the filters. In earlier studies (Schwab et al., 1996; Schwab, 1997), I limited the number of iterations, and I enlarged and pre-whitened the training patches. However, the prediction-error maps were still white.

The data dependent attribute measures – misfit of local plane wave and the prediction error – depend on local estimation within stationary local image patches. Since patches that contain a discontinuity are nonstationary, the attributes usually yield a large output anywhere within such a patch. (Figures 26 or 36. Note, that the image is composed of about 300 patches. A side of the cubic patch is 12 pixels long; about one seventh of the total data cube edge of 80 pixels.) The schemes detect if a patch contains a discontinuity, but do not pinpoint the discontinuity within the patch. Consequently, the data dependent schemes are limited to a resolution the size of their patches.

CAUSES OF SEISMIC IMAGE DISCONTINUITY

Sedimentary rocks, the standard environment of hydrocarbon reservoirs, are deposited in horizontal, parallel layers. Subsequent subsidence, uplifting or thrusting may displace or tilt a sedimentary rock package regionally; however, locally, the sedimentary rocks remain layered. Often sedimentary rocks even remain horizontal, albeit severely dipping beds are not uncommon. The detail of a seismic image volume in Figure 2 shows the typical layers of a seismic image of sedimentary rock packages. Figure 3 stylizes the layering in a dipping plane-wave volume.

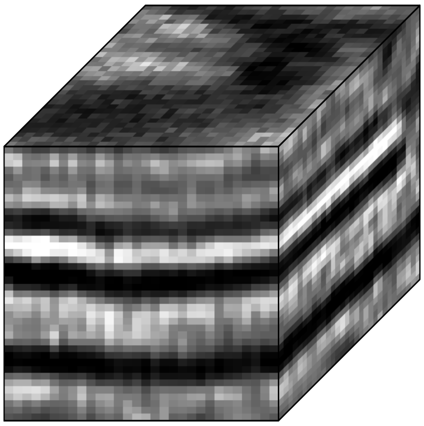

Figure 2: Detail of a typical seismic image volume. The image of the sedimentary bedding approximates a nearly horizontal plane-wave volume. nseaFoltZomRaw [ER]

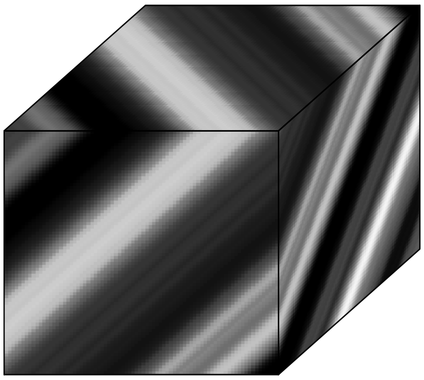

Figure 3: Synthetic plane-wave image volume. The plane-wave volume resembles undistorted sedimentary beds. planWaveRaw [ER]

In basins (the standard object of seismic images), faulting, intrusion, and erosion break the original sedimentary layers into smaller sedimentary packages and give rise to plane-wave discontinuities within the seismic image volume. Rocks of a volcanic intrusion, a turbidite flow, or a sedimentary salt give rise to discontinuities by placing their unbedded rocks next to the older sedimentary rocks. Erosion creates subsurface discontinuities by first removing parts of the original sediments and later replacing them with younger sediments. In this case, the discontinuity is due to the difference in the dip and sequence of layers of the old and young sediments. In particular, river channels wash old sediments away until it replaces them by younger layers of silt. Similarly, a discontinuity due to a fault does not separate fundamentally different rocks. Instead, a fault simply offsets and tilts the previously continuous sedimentary layers. A fault perfectly parallel to its surrounding sedimentary layers, or a fault that perfectly realigns a periodic layer package by a multiple period, would be locally invisible.

In a seismic image volume, a subsurface plane-wave discontinuity usually gives rise to a two-dimensional reflection surface, since in general the discontinuity is a boundary between rocks of different impedance. If the adjacent rocks are similar, as it is often the case with faults, the alignment of the original layers can cause reflections of variable amplitude along the discontinuity. Due to the spatial limits of a seismic survey, reflections from near-vertical discontinuities may not be recorded. Such discontinuities are only visible due to the contrast of the adjacent rock packages. For example, a fault without reflection is visible due to the misalignment of the adjacent sedimentary layer packages.

Additionally, shortcomings in the seismic experiment and processing can cause unwanted image discontinuities. Dead traces, acquisition footprints, unfocused diffractions, and large increments between time slices can result in discontinuities of the seismic image volume that are unrelated to the underlying geology. Image processing can cause otherwise sharp discontinuities to appear smooth. It can either fail to focus the discontinuities or simply smoothen it due to filtering.

In summary, I expect a discontinuity attribute to

- delineate boundaries between neighboring sedimentary packages, such as faults and large channel sands,

- indicate boundaries between sedimentary rocks and rocks without continuous reflections, such as salt and volcanic rocks,
- detect experimental or processing problems, such as acquisition footprints or unfocused diffractions.

Discontinuities vary in their sharpness and contrast due to their origin or the image's processing. In particular, reflectionless faults are sharp discontinuities and faults with reflections are comparatively smooth discontinuities.

QUALITY OF DISCONTINUITY ATTRIBUTES

The geological interpretation of seismic image volumes depends greatly on the careful mapping of its faults. The exact location and throw of faults are an important ingredient in a geologist's reconstruction of the subsurface's historic sedimentary and structural processes. Additionally, faults potentially trap or leak hydrocarbons. Picking fault surfaces in seismic image volumes is time consuming and error prone, since a fault is easily overlooked. Additionally, the fault lines, which are picked in several two-dimensional image sections, have to be combined into fault surfaces of the image volume.

Current discontinuity attributes cannot replace the picking of faults by interpreters, but they offer an alternative view and they are quick, reproducible, and three-dimensional.

The quality of a discontinuity map ultimately is the ease with which it is interpreted reliably. To that end, a discontinuity attribute needs to suppress the image's sedimentary bedding elements and isolate the image's faults and other discontinuities. Furthermore, the image should resolve any discontinuity as well as possible. In general, the discontinuity image should be free of noise and high in contrast. Beyond such fundamental criteria, every interpreter may have personal preferences.

I found it difficult to develop an image processing algorithm without the guiding principles of physical laws.

THREE TEST CASES

Synthetic image

To illustrate individual discontinuity attribute computations, I will apply them to one synthetic and two seismic image volumes. The synthetic image of Figure 4 simulates a simple fault without fault reflection. The image cube shows two distinct regions. One region is filled with a horizontal plane wave, the other with a dipping plane wave. The contact plane between the two regions models a subsurface fault. An ideal corresponding discontinuity attribute map is zero everywhere but along the fault plane.

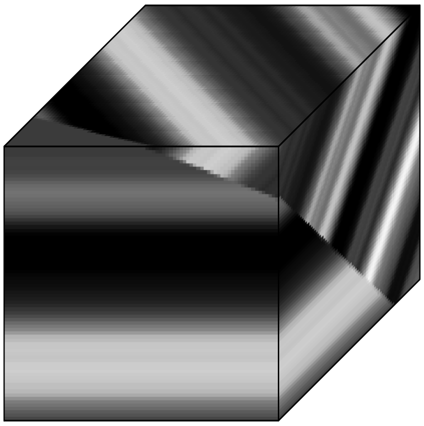

Figure 4: Synthetic fault image volume. Two sedimentary layer packages of different dip are separated by a fault without fault reflection. This image serves as an idealized test case for this chapter's discontinuity attributes. zeroFoltRaw [ER]

First seismic image

Figure 5 introduces the first seismic test case: A Gulf salt dome and its surrounding sediments and faults. The Figure 7 depicts the time slice and vertical sections indicated by the thin mid-volume lines of Figure 5.

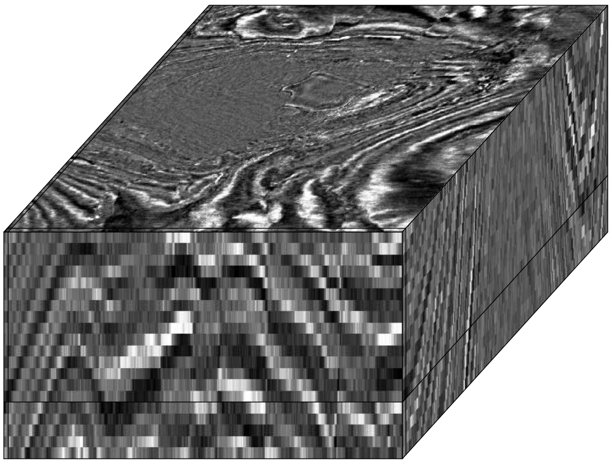

Figure 5: Volume view of Gulf salt dome. The Figure represents a volumetric view of the sections of the Figure 7. The mid-volume lines indicate the position of the extracted and displayed two-dimensional sections. gulfFoltTotRawVol [ER]

The central salt body is characterized by its low-amplitude internal reflections. The wave-like character of the time slice is due to the outcrops of steeply dipping sedimentary layers. The vertical east-west section of Figure 7 shows how the central salt body has pushed the surrounding sedimentary layers to a near vertical dip. Layers further away from the salt remained approximately horizontal.

The time slice shows a major normal fault that strikes from the north-east to the south-west and that truncates the southern tip of the salt body. Other, smaller faults simply radiate outward from the central salt dome. In addition, the image seems to be somewhat corrupted by weak north-south streaks at the nearly horizontal layers of the time slice's south-east corner. The streaks are probably due to acquisition footprint and constitute image discontinuities as well.

The vertical sections of Figure 7 show that the faults dip almost vertically and cut the sediments sharply. The left panel of Figure 6 isolates a typical fault of the image. The corresponding averaged trace spectrum on the left indicates that the fault shows considerable energy at high, near-Nyquist frequencies. In this regard, the salt dome's discontinuities resemble the synthetic test case of Figure 4.

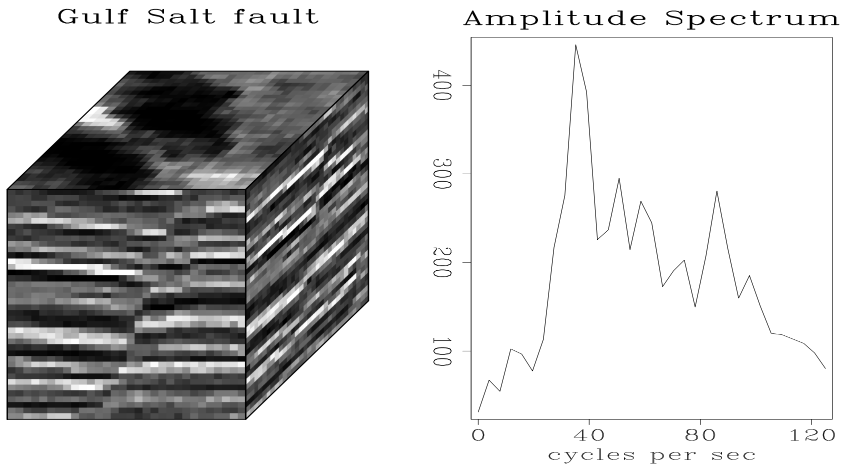

Figure 6: Gulf salt dome fault and its trace spectrum. The image detail on the left isolates a typical fault of the Gulf salt dome image. The averaged trace spectrum on the right illustrates the high-frequency contents of the fault image. $\boxed{\text{gulfFoltFltRaw}}$ [ER]

Bahorich and Farmer (1995) published discontinuity maps computed from the same seismic image volume. To be competitive with these published maps, a discontinuity attribute has to delineate the salt dome's outward radiating faults and suppress the image's sedimentary layers independent of their dip. For example, I expect a discontinuity attribute to isolate the forked fault that I marked with an `R` in the time slice and the in-line section of Figure 8. Ideally, a discontinuity map would additionally indicate the fault feature that truncates the salt body at `T`, and the pentagonoid image feature `S`.

Second seismic image

The second seismic test case, shown in Figure 9 depicts a fault-bounded North Sea horst structure. Figure 11 shows a plane view of the the selected time and vertical image slices.

The vertical sections show sets of parallel listric faults that flank a complex horst structure. Except at the northern edge of the image, the sediments are nearly horizontal. The lower fault blocks are, however, slightly tilted, which undulates the upper sediments. In contrast to the vertical section's listric faults, the time slices' intricate zigzag and rhomboidal fault pattern is difficult to discern and indicates the three-dimensional complexity of the entire horst region. The vertical sections of Figure 11 show listric faults: their dip decreases with depth. However, the faults do not sharply cut the surrounding sediments but, rather, seem to interpolate between them. As illustrated in Figure 10, the images faults show a width of a few samples and their spatial frequency components do not show much energy above half-Nyquist.

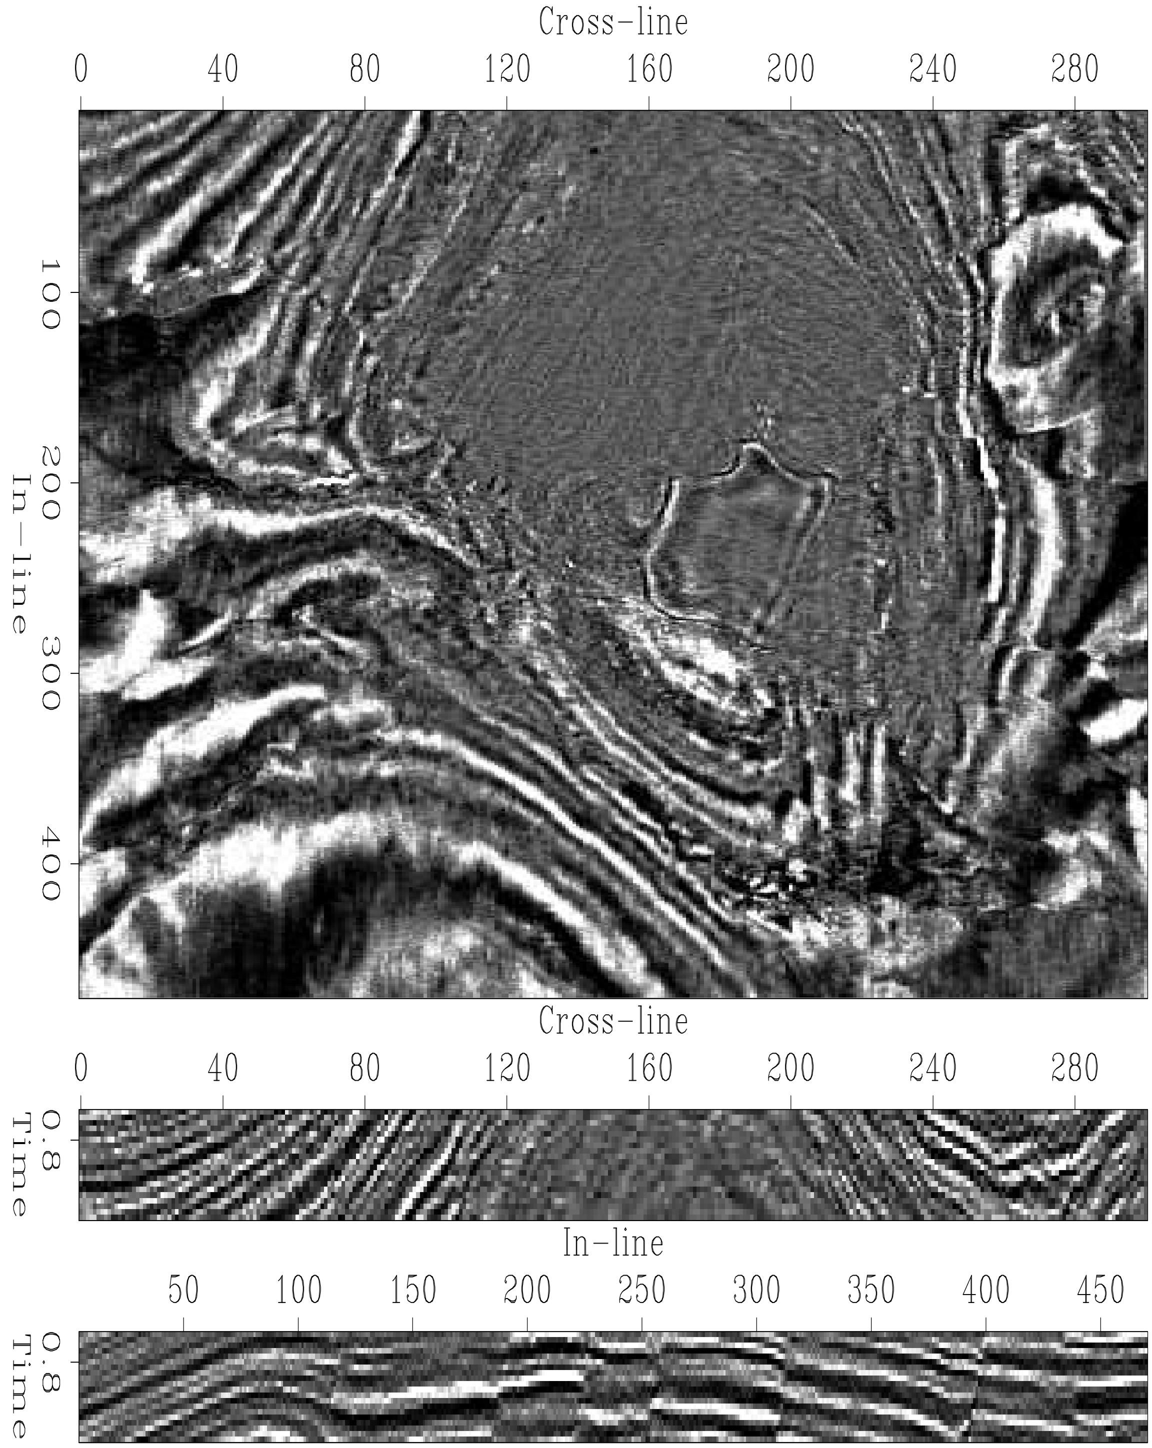

Figure 7: Gulf salt dome image. The image depicts a salt body, its surrounding sedimentary layers, and faults. gulfFoltTotRaw [ER]

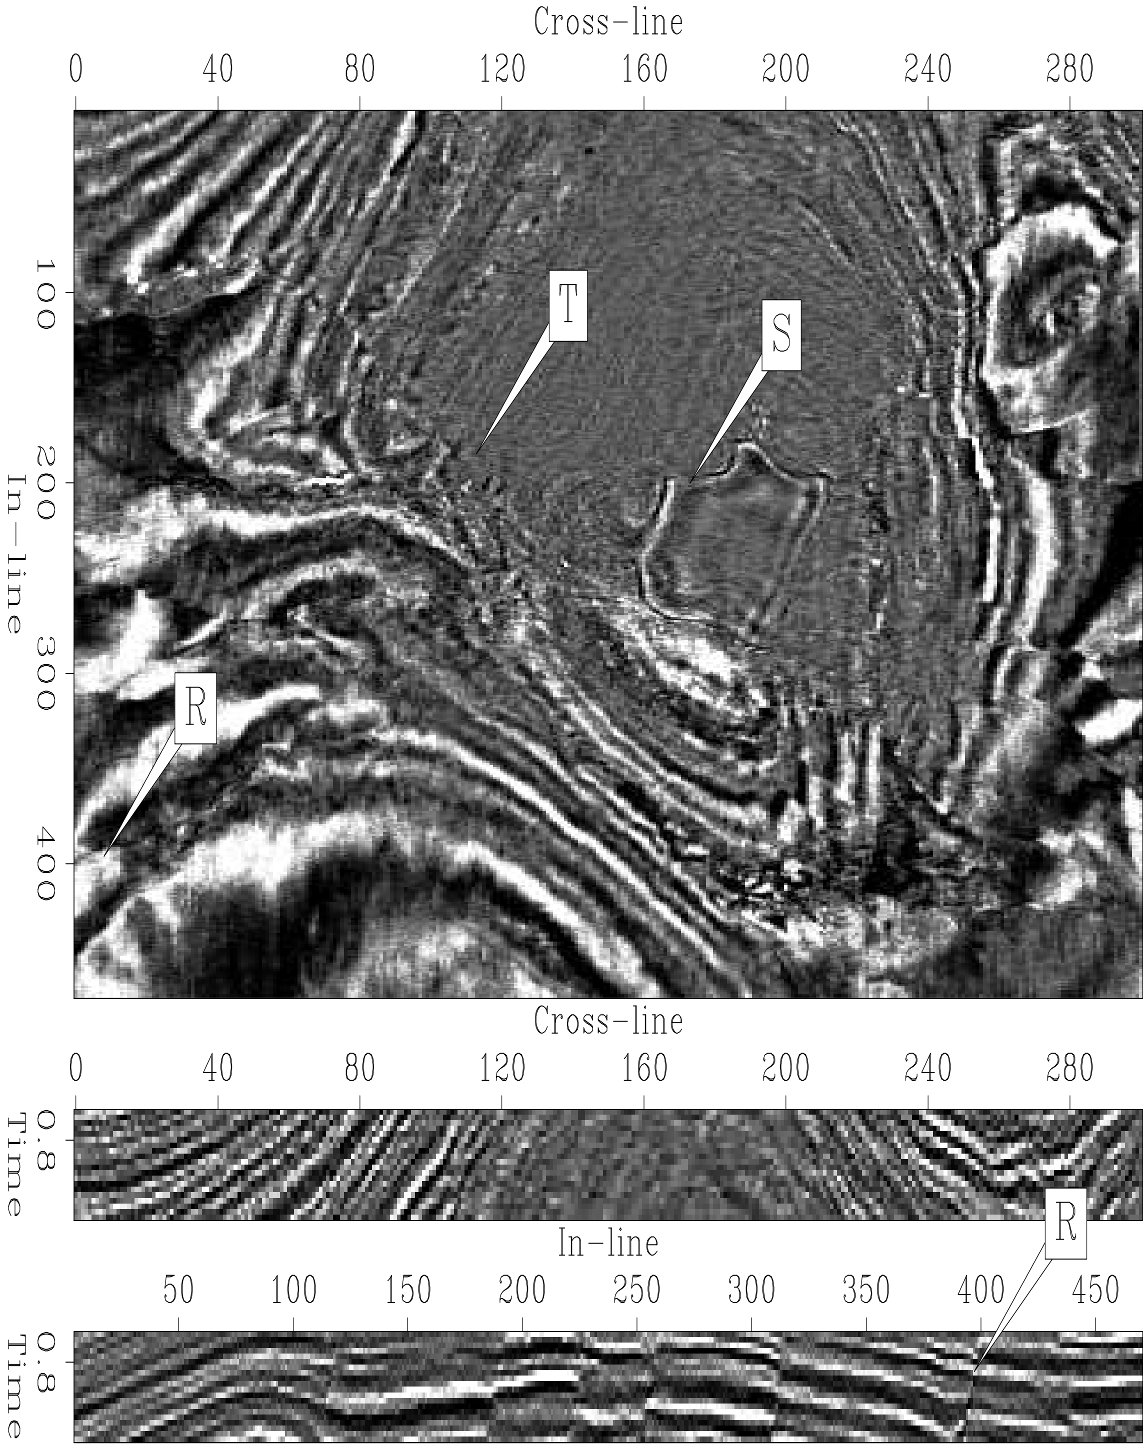

Figure 8: Annotated Gulf salt dome image. R marks a radiating fault. T indicates the salt truncation. S points at the salt body's central pentagonoid feature. gulfFoltTotMrk [ER]

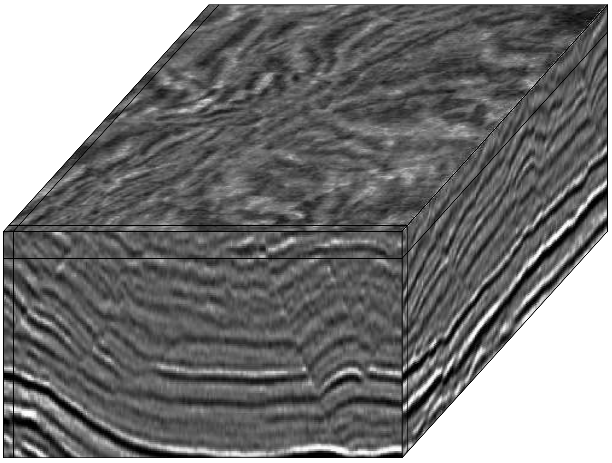

Figure 9: Volume view of North Sea horst. The Figure represents a volumetric view of the sections of the previous Figure 11. The mid-volume lines indicate the position of the extracted and displayed two-dimensional sections. nseaFoltTotRawVol [ER]

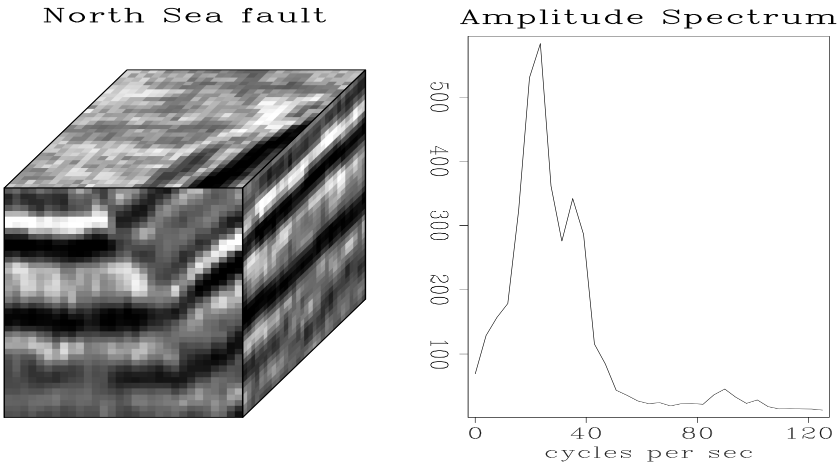

Figure 10: North Sea horst fault and its trace spectrum. The image detail on the left isolates a typical listric fault of the North Sea horst image. The averaged trace spectrum on the right demonstrates the rather low-frequency contents of the fault image. nseaFoltFltRaw [ER]

I expect a successful discontinuity map to delineate the complex fault pattern of the horst's time slice. Furthermore, I expect the sedimentary layers in the vertical sections to be removed. Bednar (**?**) tested his discontinuity attribute at this image. To facilitate a comparison among this chapter's discontinuity attributes, I marked a major fault in Figure 11 with an `F`.

DISCUSSION

Seismic volumes are difficult to interpret, since the ever present sedimentary bedding and a wide variety of other geological features usually crowd the seismic subsurface image. The image is opaque and is usually inspected in extracted two-dimensional sections.

The human eye excels in the recognition of subtle trends in otherwise noisy and incoherent data. On the other hand, the eye being restricted to two-dimensional sections is handicapped in grasping the image's volumetric features. For example, within a two-dimensional section, a fault is invisible if it intersects a sedimentary layer in strike direction. If the eye could assimilate the opaque image volume, the presence of the fault would easily be deduced from its existence in nearby image regions. However, the identification of geological features in seismic sections by an interpreter is time consuming and subjective (nonreproducible). Discontinuity maps cannot replace the human interpretation of seismic sections. The maps can, however, help to locate faults, river channels, and other geological boundaries

The synthetic seismic data set represents an idealized sedimentary fault. The fault is simply defined by the offset and of the adjacent plane-waves volumes. In that sense, the synthetic does not represent a fault reflection or transition zone of gnarled rock.

In the two seismic test cases, the major discontinuities are faults. The Gulf salt dome image additionally contains major salt boundaries. With respect to discontinuity attributes, the most important difference between the two seismic examples is the sharpness of their faults. The fault surfaces of the Gulf salt dome image 7 are significantly sharper than the fault surfaces of the North Sea horst image 11.

The image volumes of the two seismic test cases differ. The size difference of the two image volumes slightly complicates a direct comparison. The larger salt dome time slice appears less grainy than the horst's. Furthermore, little blemishes are bound to disappear in the larger image. Also, the vertical sections of the horst display more than double the time of the corresponding salt sections. Naturally, the images are much more illuminating when displayed on the screen of a graphics workstation.

ACKNOWLEDGMENTS

By its heuristic nature, a study of discontinuity attributes without field data is futile. Thanks to the companies – Haliburton, Geco-Prakla, and Unocal – that supplied migrated image volumes to test the various attribute computations. I especially want to thank Richard Day

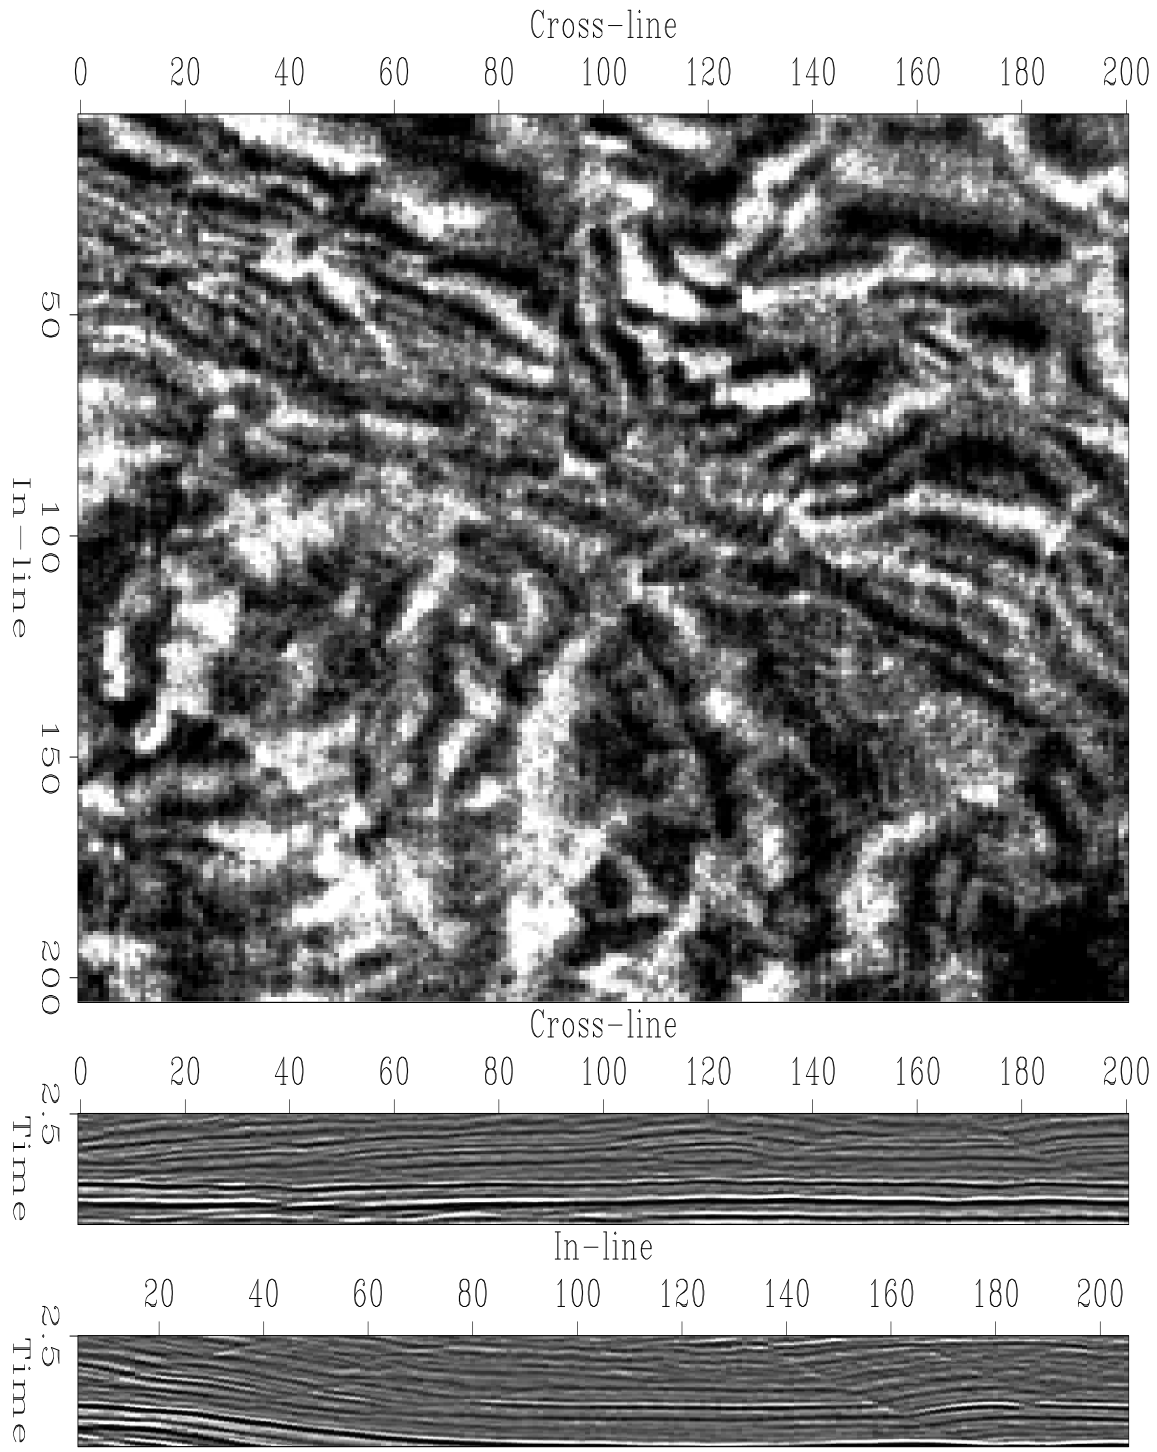

Figure 11: North Sea horst image. The time slice depicts a fault-bounded horst structure and indicates an intricate fault pattern. nseaFoltTotRaw [ER]

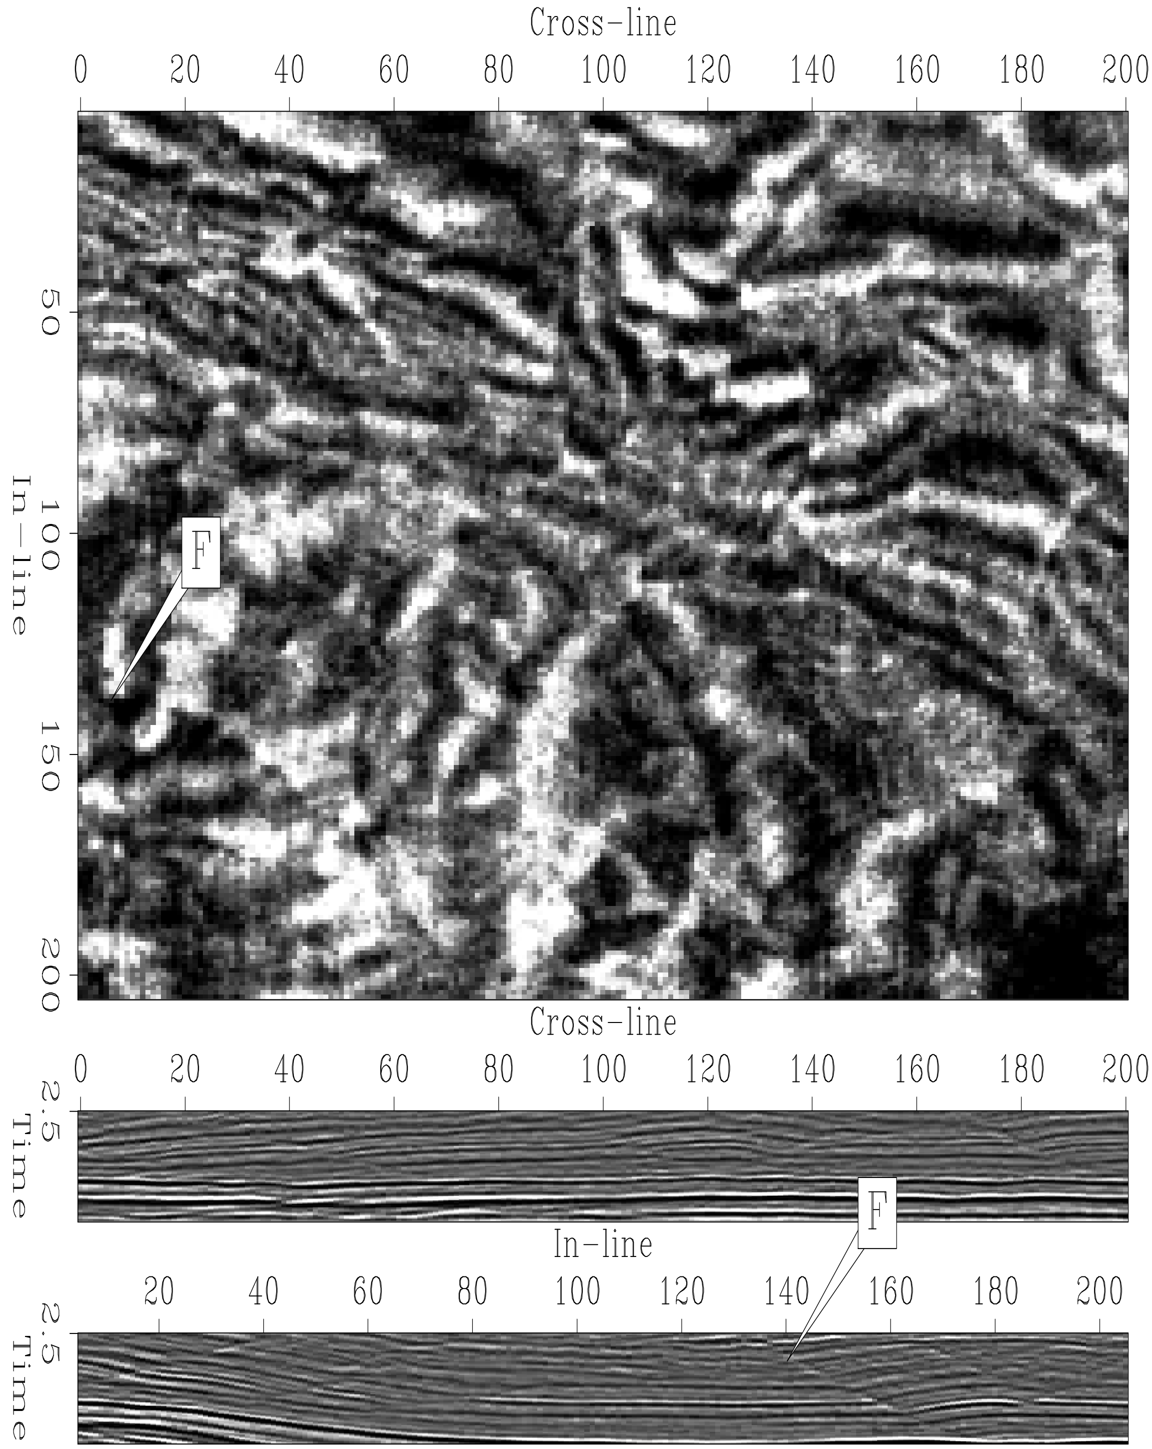

Figure 12: Annotated North Sea horst image. `F` marks a listric normal fault that is barely visible in the time slice, but that is well visible in the in-line section. nseaFoltTotMrk [ER]

of Conoco, Kurt Marfurt of Amoco, and Bee Bednar of Bee Consulting for their help in making some of these image volumes available.

Standard edge detection

Standard edge detection, such as the gradient magnitude or the Laplace operator, amplifies image regions of large local amplitude change and suppresses regions of constant amplitude. This chapter will demonstrate that standard amplitude edge detection unfortunately does not distinguish between rapid amplitude changes due to layers within a sedimentary package and amplitude changes due to package boundaries, such as faults. Consequently, standard edge enhancement techniques fail to delineate seismic image discontinuities among sedimentary layers.

GRADIENT MAGNITUDE OPERATOR

In Figure 13, a gradient magnitude operator detects the amplitude edges at which pixels change their gray-level suddenly. For an image volume $f(\mathbf{x})$, the magnitude of the gradient vector

$$|\nabla f| = \sqrt{\left(\frac{\partial f}{\partial x}\right)^2 + \left(\frac{\partial f}{\partial y}\right)^2 + \left(\frac{\partial f}{\partial z}\right)^2} \tag{1}$$

assumes a local maximum at an amplitude edge. The magnitude is zero, if f is constant.

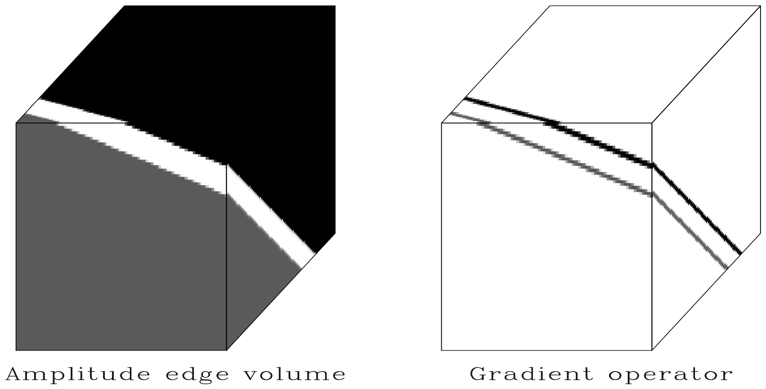

Figure 13: Gradient magnitude operator applied to constant amplitude image. The input image on the left consists of three regions of constant amplitude. The output of the gradient magnitude operator on the right delineates the edges: two parallel plane surfaces. The pixel amplitudes are all non-negative. Large amplitudes in the output image indicate edges in the input image. Zero output amplitude indicates locally constant input. $\boxed{\text{consFoltGrad}}$ [ER]

In the case of digital images, partial derivatives are conveniently computed by finite-

Roberts (?)	Smoothed Prewitt (?)	Expanded Prewitt	Sobel (?)	Isotropic
$\begin{matrix} 0 & 1 \\ -1 & 0 \end{matrix}$	$\begin{matrix} -1 & 0 & 1 \\ -1 & 0 & 1 \\ -1 & 0 & 1 \end{matrix}$	$\begin{matrix} -1 & -1 & 0 & 1 & 1 \\ -1 & -1 & 0 & 1 & 1 \\ -1 & -1 & 0 & 1 & 1 \\ -1 & -1 & 0 & 1 & 1 \\ -1 & -1 & 0 & 1 & 1 \end{matrix}$	$\begin{matrix} -1 & 0 & 1 \\ -2 & 0 & 2 \\ -1 & 0 & 1 \end{matrix}$	$\begin{matrix} -1 & 0 & 1 \\ -\sqrt{2} & 0 & \sqrt{2} \\ -1 & 0 & 1 \end{matrix}$

Table 2: Finite-difference approximations of the horizontal partial derivative.

difference approximations. In particular, I use a three-dimensional generalization of

$$\frac{\partial}{\partial x} f(i, j) = f(i, j) * m_x(i, j) \text{ where } m_x = \begin{bmatrix} 1 & 1 \\ -1 & -1 \end{bmatrix} \tag{2}$$

to approximate the horizontal partial derivative of this chapter's image volumes. Table 2 lists alternative finite-difference approximations of the partial derivative operator. The various masks vary in their moments, which affects the operators' isotropy, robustness and resolution. Averaging of surrounding values before subtraction reduces the effect of noise and increases the sensitivity to gradual grey-level changes. However, an increase in operator size causes a decrease in the operator's ability to resolve two neighboring edges. The finite-difference approximations illustrate the prediction character of the filter: the difference between the two neighboring pixels is the error one makes predicting one amplitude value to be the same as the next.

The Derivative Theorem of the Fourier Transform,

$$\begin{array}{lrcl} \text{If} & \mathcal{F}[f(x, y)] & = & F(u, v) \\ \text{then} & \mathcal{F}[\frac{\partial}{\partial x} f(x, y)] & = & i 2\pi u\, F(u, v), \end{array}$$

shows that the partial derivative operator scales each Fourier component of the input image by its spatial wavenumber in the derivative direction. Consequently, the high frequency components that constitute an edge are amplified and low frequency components that constitute near constant amplitude areas are suppressed. The zero frequency component is erased. Additionally, the derivative operator causes a 90^0 degree phase shift.

Results

The gradient magnitude operator does have its desired edge enhancing effect on non-oscillatory images, such as the synthetic data in Figure 13 or the Bay Area map in Figure 14. The synthetic test image on the left of Figure 13 consists of three constant-amplitude regions that are separated by two parallel planes. The gradient magnitude map on the right delineates the two boundary planes.

The top panel of Figure 14 shows the Bay Area topography. The gradient magnitude operator distinguishes topographic features by their slope. In contrast to the Bay Area's smooth sloping hills, the magnitude map's multiple areas of zero gradient surprisingly imply a step-like landscape. I assume that the original topography data was probably sloppily interpolated from a contour map. In general, edge detection and discontinuity attributes often reveal such image shortcomings.

However, applied to oscillatory images, such as the synthetic test case in Figure 15, the gradient magnitude operator indiscriminately amplifies the plane-wave layer boundaries as well as the discontinuity between the plane-wave volumes. Consequently, the plane-wave layers are not suppressed and the central fault is not isolated. The amplitude along the enhanced discontinuity surface varies irregularly depending on the amplitude contrast of the adjacent sedimentary layers.

Similarly, the gradient magnitude operator fails to delineate the faults of the seismic image example of Figure 16. The operator increases the overall frequency contents of the image, but does not suppress the sedimentary layer packages. The steeply inclined layers appear aliased in the vertical sections. The faults are hidden among the layer boundaries. Faults, such as the exemplary `R` fault of the original image 8, can be identified only if one knows where to search for them. The so-identified faults are, however, often only one-pixel wide and, consequently, well resolved. The salt truncating fault is not particularly enhanced. Its position is indirectly indicated by a change in the salt bodies average gradient magnitude. Similarly, the boundaries of the salt dome are not clearly pinpointed.

In the case of the Gulf salt dome and the North Sea horst, the gradient operator fails to isolate the faults and yields a map that is as difficult to interpret as the original seismic image.

IMAGE AND PHYSICAL SPACE

In this chapter, I compute discontinuities of a seismic image, not discontinuities of the depicted subsurface. The computations are based on the set of image pixels and do not directly relate to the physical subsurface that the pixels represent. Consequently, partial derivatives indicate change per sample and not change per physical unit, such as meters or seconds. Two differently sampled images of a single physical field $f(x, y, z)$ will result in two different discontinuity images.

The pixel space is defined by the uniform samples of a given discrete image. Its variables $(\hat{x}, \hat{y}, \hat{z})$ count samples. The vector space has the natural norm $\sqrt{\hat{x}^2 + \hat{y}^2 + \hat{z}^2}$. Convolutions such as expression (2) conveniently compute approximations of the partial derivatives of the image. All derivative expressions share the unit *amplitude change over sample*. The gradient magnitude expression (1) combines the partial derivatives of the sample space.

A change of variables relates the pixel coordinates $(\hat{x}, \hat{y}, \hat{z})$ to the underlying physical coordinates (x, y, z):

$$\begin{aligned} x &= x_0 + \hat{x}\, dx \\ y &= y_0 + \hat{y}\, dy \\ z &= z_0 + \hat{z}\, dz, \end{aligned} \tag{3}$$

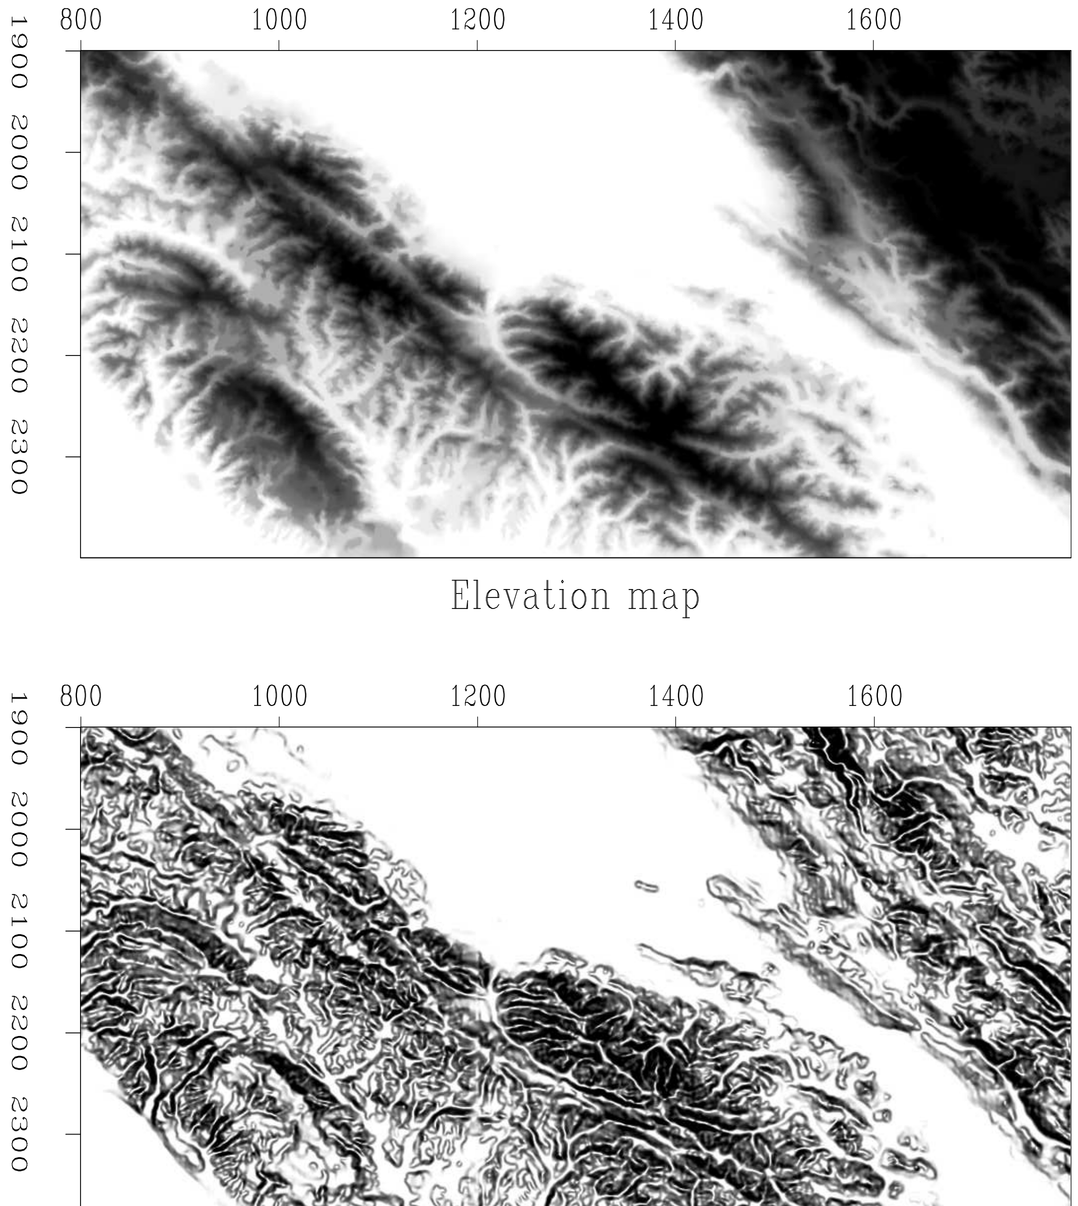

Figure 14: Gradient magnitude operator applied to Bay Area topography. The top panel shows the topographic elevation of the San Francisco Bay Area. In the bottom panel, the gradient magnitude operator successfully enhanced the topographic *edges*, such as ravines and canyons. The terraced appearance is an artefact of the original data. bayGrad [ER]

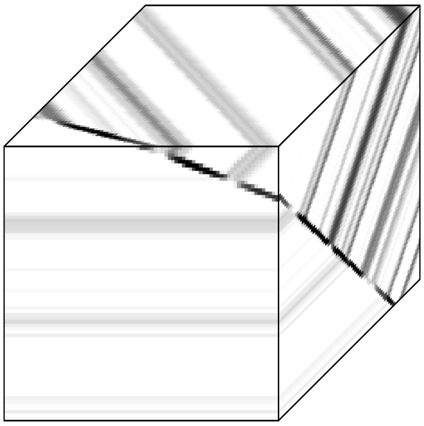

Figure 15: Gradient magnitude operator applied to synthetic test case. The gradient magnitude operator fails to isolate the fault, because it cannot differentiate between the amplitude change across the fault and the amplitude change across the sedimentary layers. zeroFoltGrad [ER]

where (x_0, y_0, z_0) are the axes offsets and (dx, dy, dz) are the axes increments between samples. The partial z-derivative of the physical space and the pixel space are related by

$$\frac{\partial f}{\partial z} = \frac{\partial \hat{f}}{\partial \hat{x}}\frac{\partial \hat{x}}{\partial z} + \frac{\partial \hat{f}}{\partial \hat{y}}\frac{\partial \hat{y}}{\partial z} + \frac{\partial \hat{f}}{\partial \hat{z}}\frac{\partial \hat{z}}{\partial z} = \frac{1}{dz}\frac{\partial \hat{f}}{\partial \hat{z}}, \tag{4}$$

where f is a physical space function and $\hat{f}$ is the corresponding pixel space function.

The physical space is more intractable than the pixel space. For example, the gradient magnitude of a time-migrated image in pixel space is well defined by expression (1). The same gradient magnitude cannot be formed in the corresponding physical space, since the units of the partial spatial derivatives and the partial temporal derivative differ.

To compute the gradient magnitude of the physical space, a time-migrated image would have to be converted from time to depth. A true time-depth conversion requires a subsurface velocity model and the combination of an inverse time migration and a depth migration. Such a costly transformation is unsuitable for a mere discontinuity attribute computation. An alternative pseudo-depth conversion, $z = v/2\,t$, is a change in variables and leads to a partial derivative weight

$$\frac{\partial f_t}{\partial t} = \frac{v}{2}\frac{\partial f_z}{\partial z}.$$

where f_t is the time-migrated and f_z is the pseudo-depth function. v is a constant approximate subsurface velocity. The expression

$$\frac{\partial \hat{f}_t}{\partial t} = \frac{v\,dt}{2\,dz}\,\frac{\partial \hat{f}_z}{\partial z}$$

links the partial z-derivative of the time migrated pixel function $\hat{f}_t$ to the partial z-derivative of the pseudo-depth pixel space $\hat{f}_z$.

Seismic images are usually oversampled in time (small dt), and barely sampled adequately in space. Consequently, the time-depth conversion factor dz/dt usually increases the weight of the pixel z-derivative. In the case of the Gulf salt dome image of Figure 7,

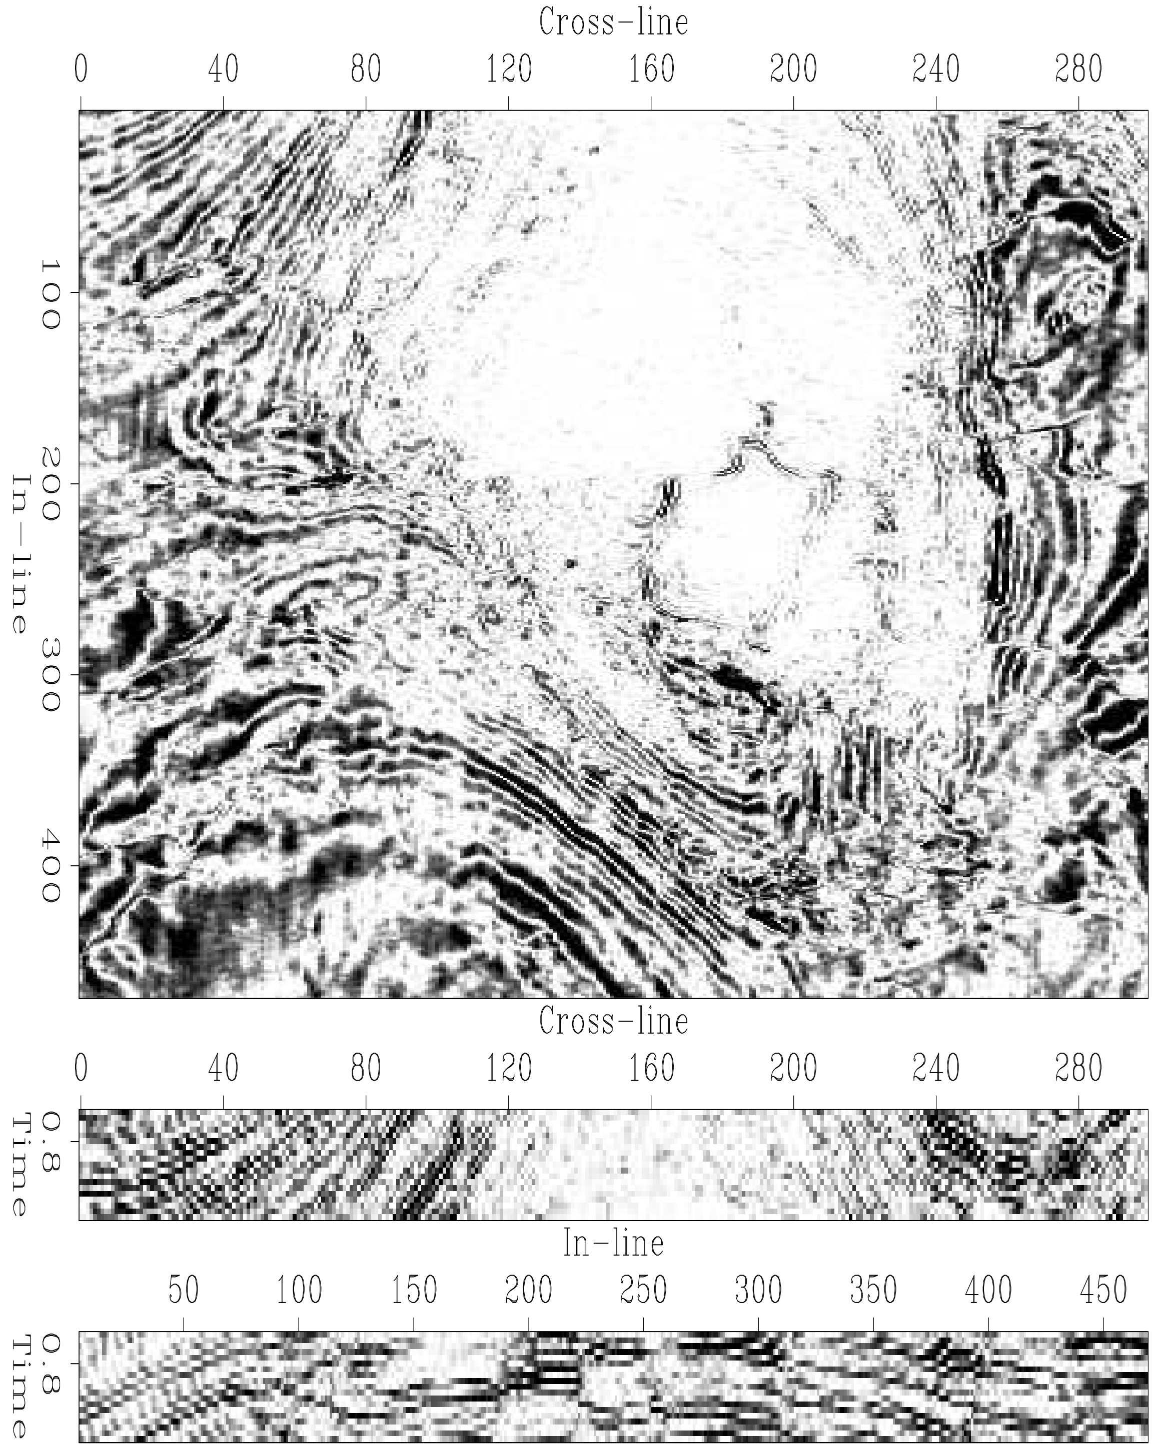

Figure 16: Gradient magnitude operator applied to salt image. The gradient magnitude operator fails to suppress the image's sedimentary layers. gulfFoltTotGrad [CR]

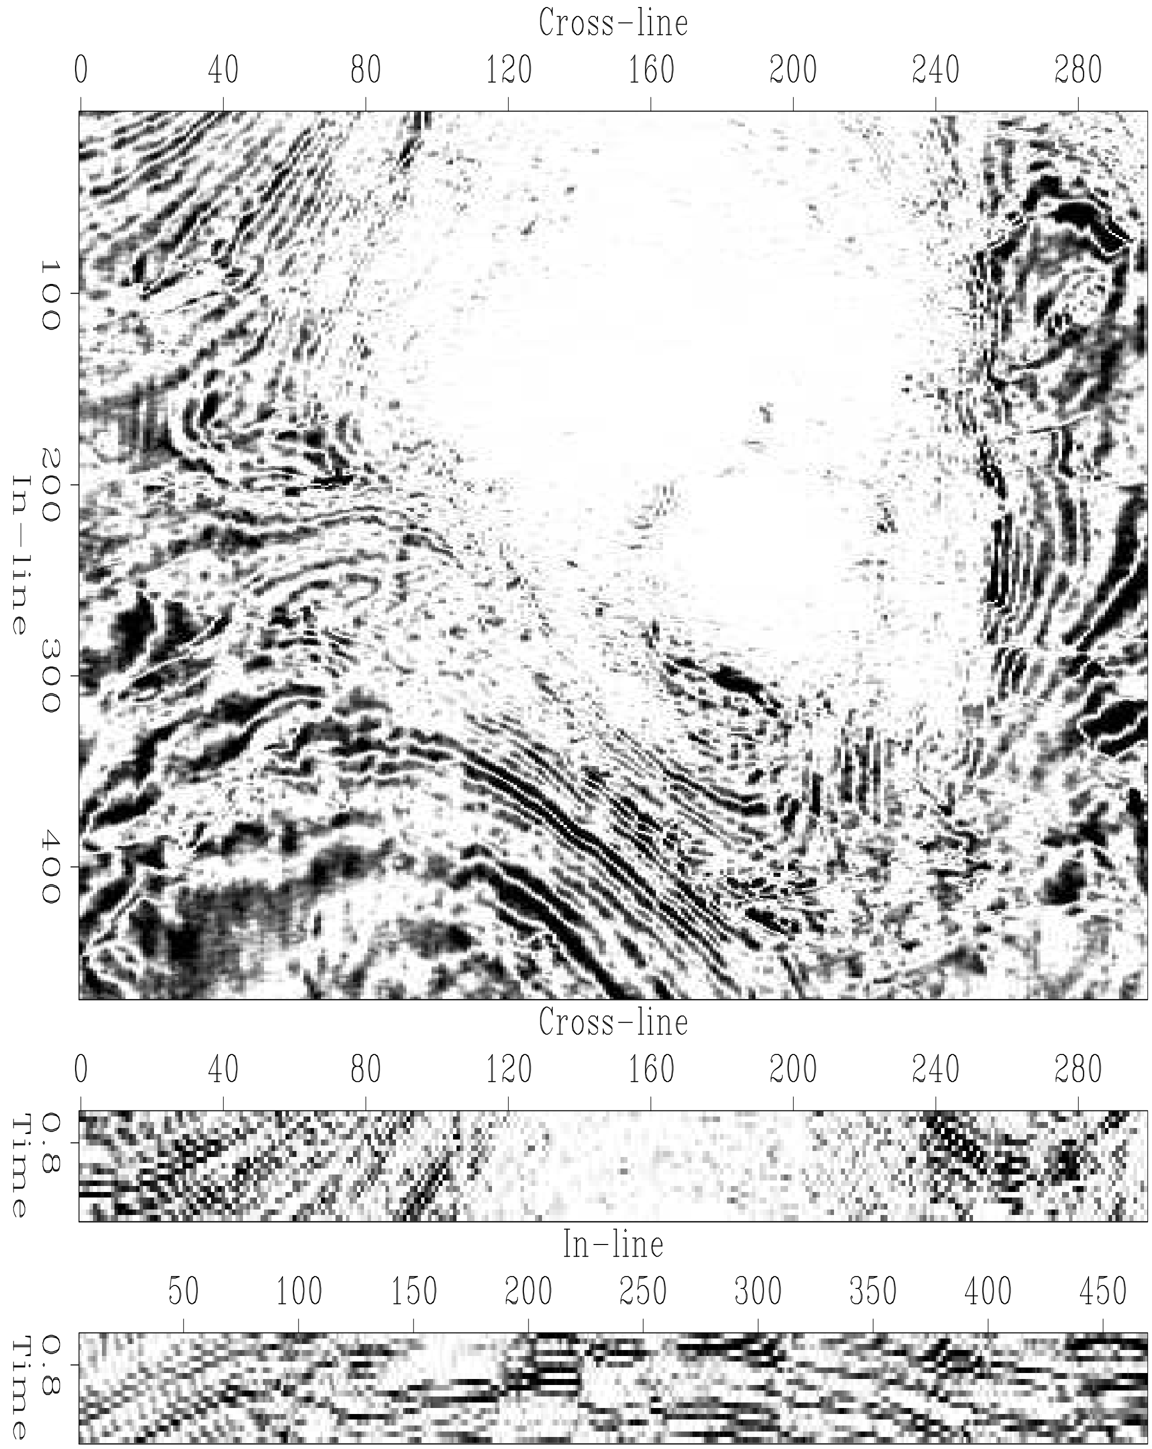

Figure 17: Weighted gradient magnitude operator applied to salt image. In contrast to Figure 16, this image is the gradient magnitude of the pseudo-depth space, not simply the pixel space. gulfFoltTotWGrd [CR]

the relation between the two partial derivatives is

$$\frac{\partial \hat{f}_z}{\partial z} \approx 8.3 \frac{\partial \hat{f}_t}{\partial t}$$

since $dz = 0.05km$, $dt = 0.004s$, $v = 3.0km/s$.

Figure 16 shows the gradient magnitude image of the first seismic test case computed in pixel time space. Figure 17 shows the corresponding magnitude image based on the pseudo-depth space. Both images show little difference, except that the salt body appears noisier in the unweighted image.

In general, I choose the *unphysical* pixel space to compute this chapter's discontinuity attributes. Why? The computation of a discontinuity attribute is not based on a physical process and it makes little sense to talk about its correctness. Its usefulness results purely from its interpretational clarity and its computational simplicity.

In my experience, a large weight of the temporal pixel derivative decreases the quality of the discontinuity map. I believe the deterioration is due to the importance of the horizontal derivative terms in annihilating the prevalent horizontal sedimentary layers. Empirically, the rather successful discontinuity attribute of Figure 23 demonstrates the importance of horizontal comparison. The attribute is horizontal correlation.

Furthermore, the pixel space measurements are conceptually and computationally simple. Traditionally, enhancement techniques in the image processing literature use the same pixel space approach (Castleman, 1996).

PROCESSING AND DISPLAY PARAMETERS

To ensure that the result images, such as Figure 7 and 16, are comparable, I process identical image volumes for all computations. All results are represented by the same image sections of the image volume.

Usually, each processing step offers its user a set of configurable parameters. In the case of the gradient magnitude operator, users can set individual weights for each partial derivative term or choose among competing finite-difference approximations of the partial derivatives. In general, I experiment with such processing parameters. Because of the shear amount of possible combinations, such experiments are usually not exhaustive. The result images that I include in this report are the *best* results that I found. Similarly, I choose individual display parameters, such as polarity and clip, for each image.

LAPLACE OPERATOR

In Figure 18, a Laplace operator transforms the amplitude edges in the image cube on the left to zero crossings of a bipolar wavelet in the image cube on the right. Given an image volume $f(\mathbf{x})$, the Laplace operator

$$\nabla^2 f(\mathbf{x}) = \frac{\partial^2 f(\mathbf{x})}{\partial x^2} + \frac{\partial^2 f(\mathbf{x})}{\partial y^2} + \frac{\partial^2 f(\mathbf{x})}{\partial z^2}$$

zeroes the amplitude where f is constant. Analogous to the gradient magnitude operator case, a Laplace operator is conveniently approximated by its finite difference approximations, such as those shown in Table 3. In general, the masks vary in their moments and consequently in their isotropy, robustness and resolving power. In this chapter, I use a three-dimensional version of the one-dimensional and two-dimensional Laplace filters of Figure 19.

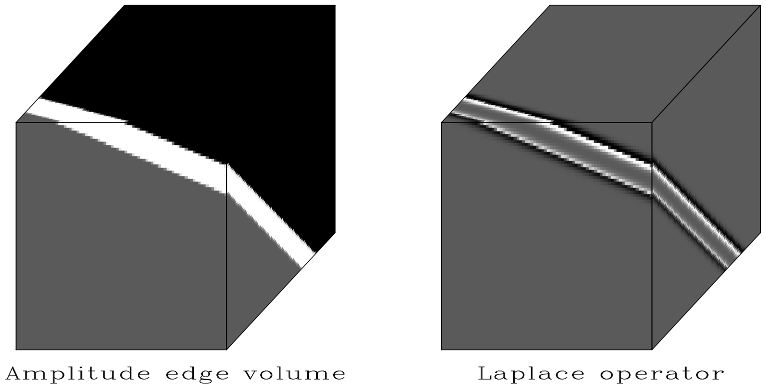

Figure 18: Laplace operator applied to constant amplitude image. The output of the Laplace operator on the right delineates the edges successfully. Edges are indicated by the zero-crossing (grey) of a bipolar wavelet. Constant input regions result in zero output values (grey). consFoltLap [ER]

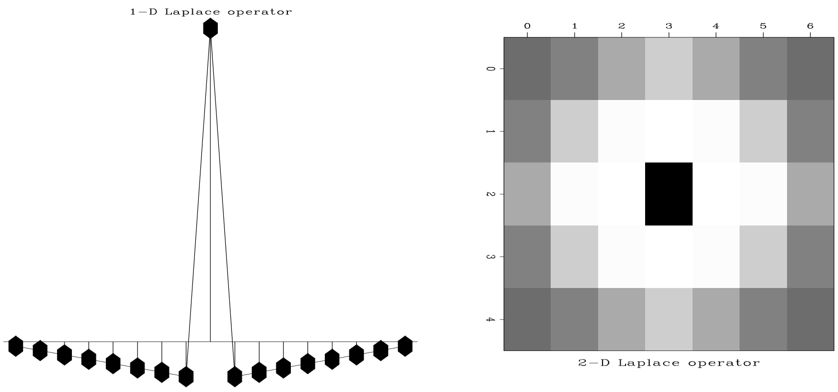

Figure 19: Generalized Laplace operator. The one-dimensional filter on the left is constructed by subtracting a wide and a narrow triangle of equal area. Similarly, the two-dimensional filter on the right is constructed by subtracting a wide and a narrow pyramid of equal volume. lapGeneral [ER]

$\begin{matrix} 0 & -1 & 0 \\ -1 & 4 & -1 \\ 0 & -1 & 0 \end{matrix}$	$\begin{matrix} -1 & -1 & -1 \\ -1 & 8 & -1 \\ -1 & -1 & -1 \end{matrix}$	$\begin{matrix} 1 & -2 & 1 \\ -2 & 4 & -2 \\ 1 & -2 & 1 \end{matrix}$	$\begin{matrix} -1 & -1 & -1 & -1 & -1 \\ -1 & -1 & -1 & -1 & -1 \\ -1 & -1 & 24 & -1 & -1 \\ -1 & -1 & -1 & -1 & -1 \\ -1 & -1 & -1 & -1 & -1 \end{matrix}$

Table 3: Finite-difference approximations of the 2-D Laplace operator.

The finite-difference representations of Table 3 approximate a generalized Laplace operator that combines the nabla operator with a pre-smoothing by a Gaussian:

$$g(x, y) = -\nabla^2 \left[\frac{1}{2\pi\sigma^2} e^{-\frac{x^2+y^2}{2\sigma^2}} \right] = \frac{1}{\pi\sigma^4} \left[1 - \frac{x^2+y^2}{2\sigma^2} \right] e^{-\frac{x^2+y^2}{2\sigma^2}}. \tag{5}$$

where σ controls the width of the Gaussian kernel. In the Fourier domain, the generalized Laplace operator (5) is

$$\mathcal{F}[g(x, y)] \propto (u^2 + v^2) e^{-2\sigma^2(u^2+v^2)}.$$

The factor $(u^2 + v^2)$ is the Fourier transform of the nabla operator; the exponential is the transform of the Gaussian low-pass filter.

Results

The Laplace operator detects edges between constant amplitude volumes (Figure 18). Applied to the synthetic test case, the Laplace operator delineates the discontinuity but only partially suppresses the plane waves (Figure 20). The wavelet along the discontinuity changes its polarity depending on the contrast of the adjacent plane waves. The somewhat lower amplitude of the Laplace operator among the horizontal plane wave is probably due to the exact zero of the horizontal component.

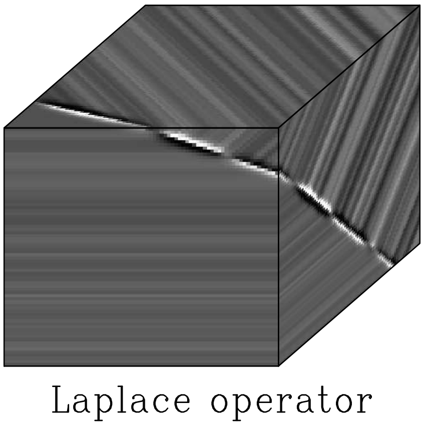

Figure 20: Laplace operator applied to synthetic test case. The Laplace operator slightly enhances the fault plane between the neighboring plane-wave volumes. However, the plane-wave oscillations lead to noticeable residuals in the plane-wave volumes. zeroFoltLap [ER]

Applied to the first seismic test case (Figure 21), the Laplace operator creates a discontinuity map that generally resembles the original seismic image. The wavefield-like character of the time slice remains. The fault marked R in the original image 8 is not enhanced: neither are any other faults. Applied to the image of the North Sea horst, the Laplace operator yields a similarly disappointing discontinuity map (not shown).

HORIZONTAL CORRELATION

In the special case, that all sedimentary layers are horizontal, a time slice of an undisturbed horizontal layer shows a constant amplitude, and, within the time slice, discontinuities create standard amplitude edges. The amplitude edges could be enhanced by standard edge enhancement techniques applied to individual time slices. Instead, I chose trace correlation to compute discontinuity attributes of horizontally layered seismic images.

Correlation measures horizontal alignment of one-dimensional arrays (traces). The standard normalized correlation between two one-dimensional arrays, f_1 and f_2, is

$$c = \frac{\sum_i f_{1i} f_{2i}}{\sqrt{\sum_i f_{1i} \sum_i f_{2i}}} \tag{6}$$

where i is the arrays' sample index. Neidell and Taner (1971) generalize the correlation coefficient to measure alignment among a set of one-dimensional arrays:

$$c = \frac{2}{N(N-1)} \sum_i \sum_{j>i} \left[\frac{\sum_k f_{ik} f_{jk}}{\sqrt{\sum_k f_{ik}^2 \sum_k f_{jk}^2}} \right], \tag{7}$$

where N is the length of a single trace in samples and f_{ik} is the k-th sample of the i-th trace. The correlation coefficients are bounded by -1 and 1. The coefficient c is 1 only if all the traces are identical.

To simplify plots of correlation coefficients, I map the coefficient range from $[-1, 1]$ to $[0, 2]$ so that pixels of correlation 1 – perfect correlation – plot as 0 and less well-correlated regions show positive amplitudes. Hence, I first subtract 1 from each pixel value and then multiply it by -1.

To compute the local correlation of a nonstationary image volume, I split the image into small patches, compute the generalized correlation within the patch, set all patch pixels to its correlation value c and merge the individual patches to a single output quilt[1]. The correlation within a patch measures the similarity among its traces f_i, the one-dimensional vertical array of all pixels of identical horizontal location. Perfectly horizontal layers result in a correlation coefficient of 1 and are mapped to a pixel value of 0.

The patch size is determined by a trade-off between resolution and data nonstationarity on one side and the need to capture sufficient statistical information in a patch on the other. Since each patch is filled with a single correlation coefficient, the procedure can only resolve two discontinuities if they are separated by more than the size of a patch. Besides loss of resolution, a large patch may capture nonstationary data and yield incorrect statistical estimates. On the other hand, a small patch may not contain enough data to gather reliable statistics in the presence of noise.

[1]This dissertation's chapter on nonstationary data processing discusses the details of the patching procedure.

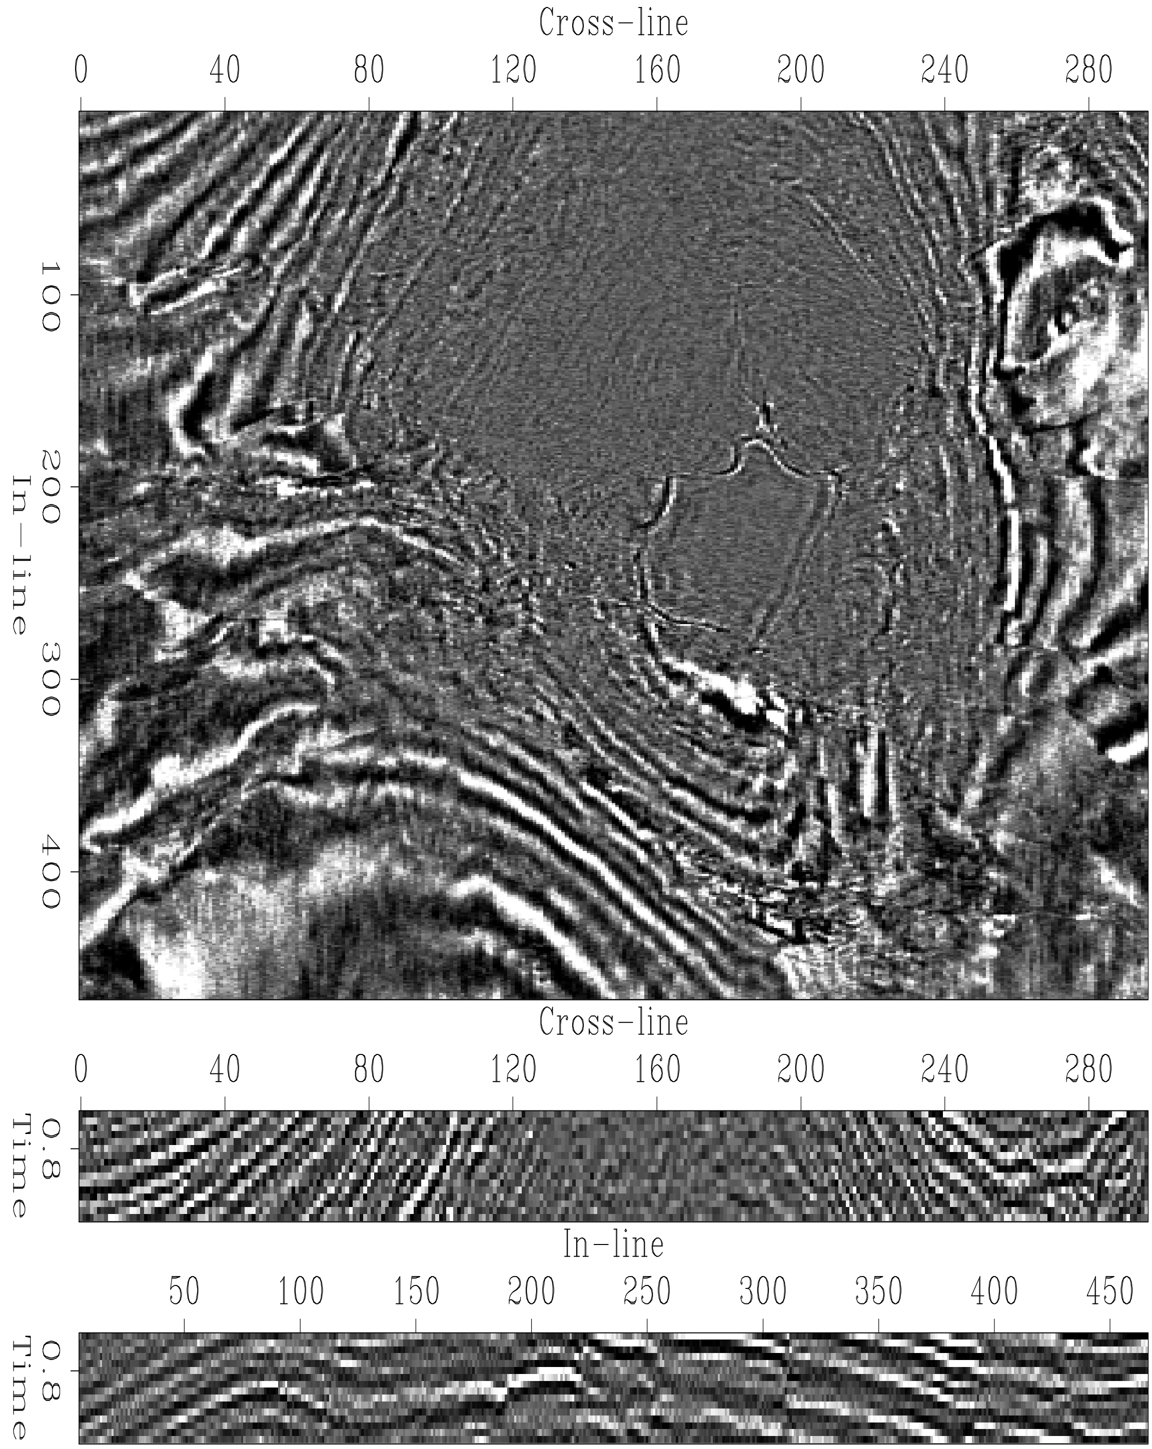

Figure 21: Laplace operator applied to salt image. The Laplace operator does not change the character of the original seismic image. In particular, the discontinuity map fails to delineate the sought faults. gulfFoltTotLap [ER]

Results

The horizontal correlation attribute suppresses the horizontal layers, but fails to remove dipping ones (Figure 22). The discontinuity is delineated where the adjacent layers sufficiently contrast. The large and blurred character of the image's features illustrates the reduced resolution due to the patch size.

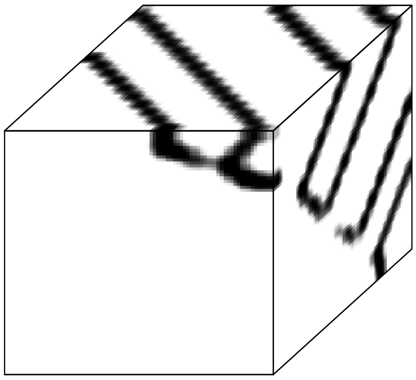

Figure 22: Horizontal correlator applied to synthetic test case. The correlation of neighboring trace intervals annihilates the horizontal beds but fails to reject the dipping ones. Hence, the fault is not well isolated. zeroFoltCor [ER]

Horizontal correlation is a surprisingly successful discontinuity attribute when applied to seismic subsurface images. Applied to the Gulf salt dome image (Figure 23), the correlation successfully suppresses the horizontal layers. Even dipping layers are removed, if laterally they smoothly vary, due to their width. However, high-frequency dipping layers, such as in the north-west corner, are not removed, but show a considerable discontinuity amplitude.

The radiating faults are discerned. The fault marked R in the original seismic image 8 is depicted. An intrepid interpreter may even suspect that the fault diagonally crosses the salt body and reappears in the north-east corner of the time slice. The primary faults in the subsurface are accompanied by similar linear features with lower amplitude. These events might be fault surfaces that only partially intersect the local patch or that separate sedimentary layers of little local contrast. Overall, the faults contrast sharply with their surrounding image regions. The discontinuity map has about patch-sized resolution, as expected.

The salt itself gives rise to a region of little correlation and consequently high amplitude. Among the chaos of events, the salt truncating faults are not enhanced. The central pentagonoid region at the southern tip of the salt body mostly vanishes in the discontinuity map.

In the case of the North Sea horst image, the horizontal correlation reveals a complex set of linear discontinuity events. The sedimentary layers are removed in the time slice and the faults of the original image can be related to the faults in the attribute image. However, the attribute fault indicates additional discontinuities. Are these discontinuities significant for the image's geological interpretation?

In Figure 24, the discontinuities contrast sharply with the surrounding sedimentary packages. However, the processes resolution is about the size of the correlation patches.

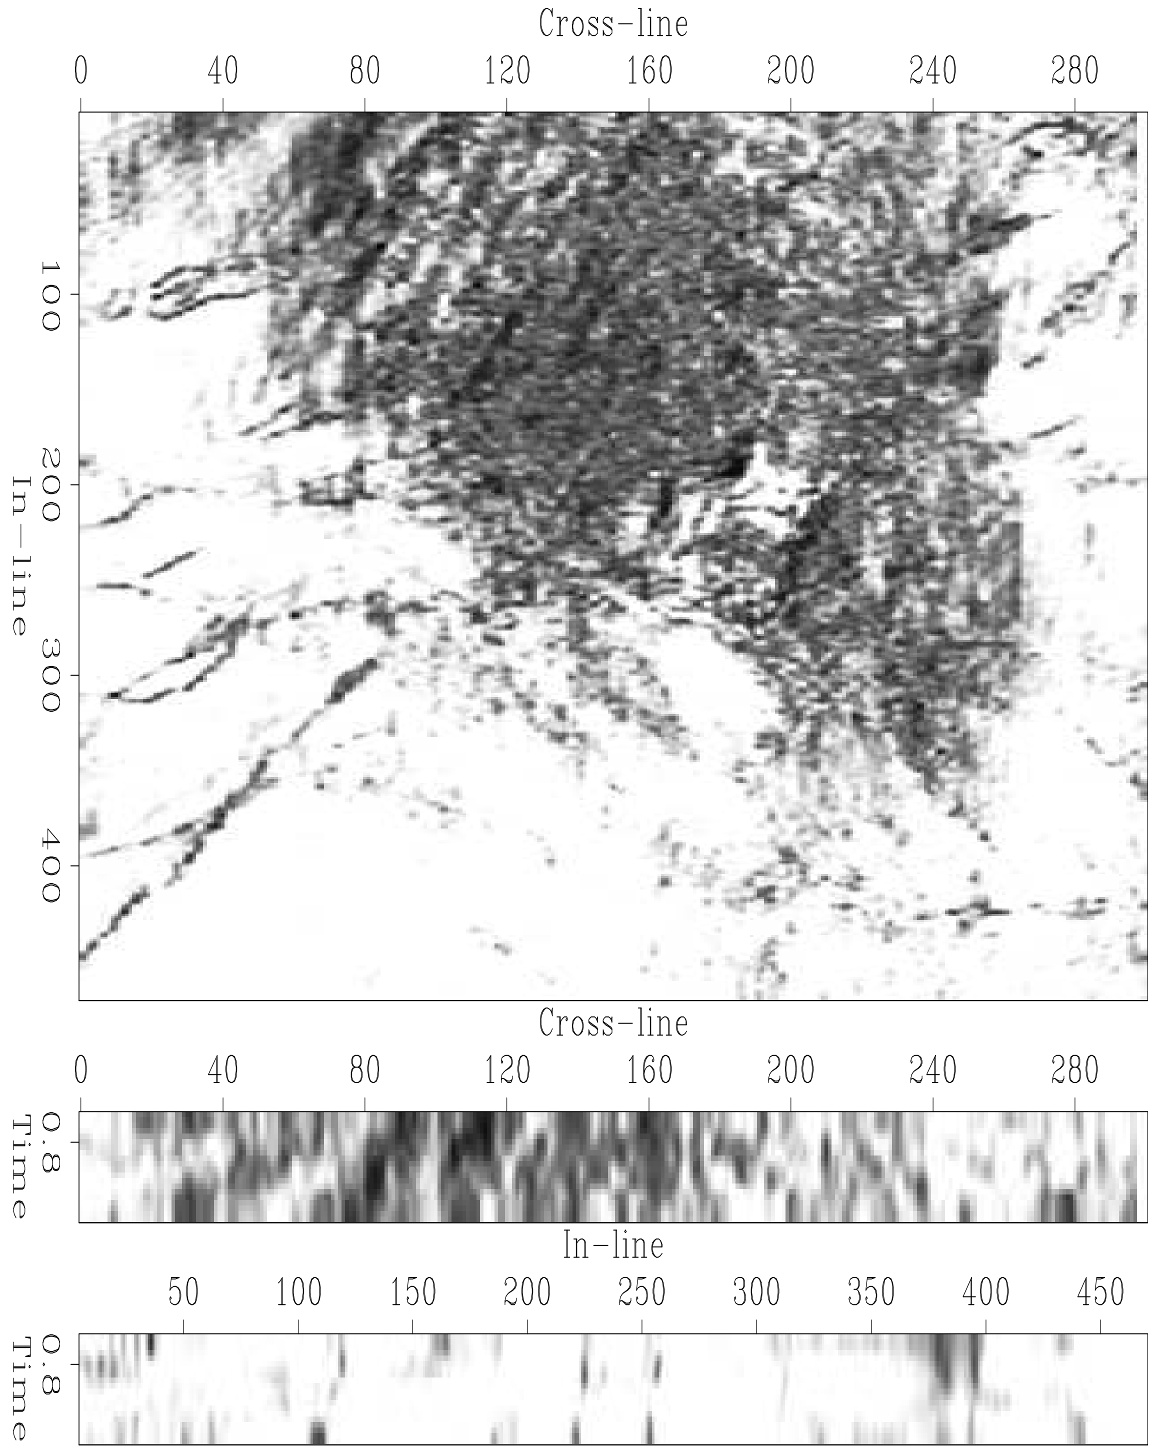

Figure 23: Horizontal correlator applied to salt image. The correlation coefficient of neighboring traces removes the image's sedimentary layers and delineates its faults. The simple method is remarkably successful, but does not resolve many details. gulfFoltTotCorr [CR]

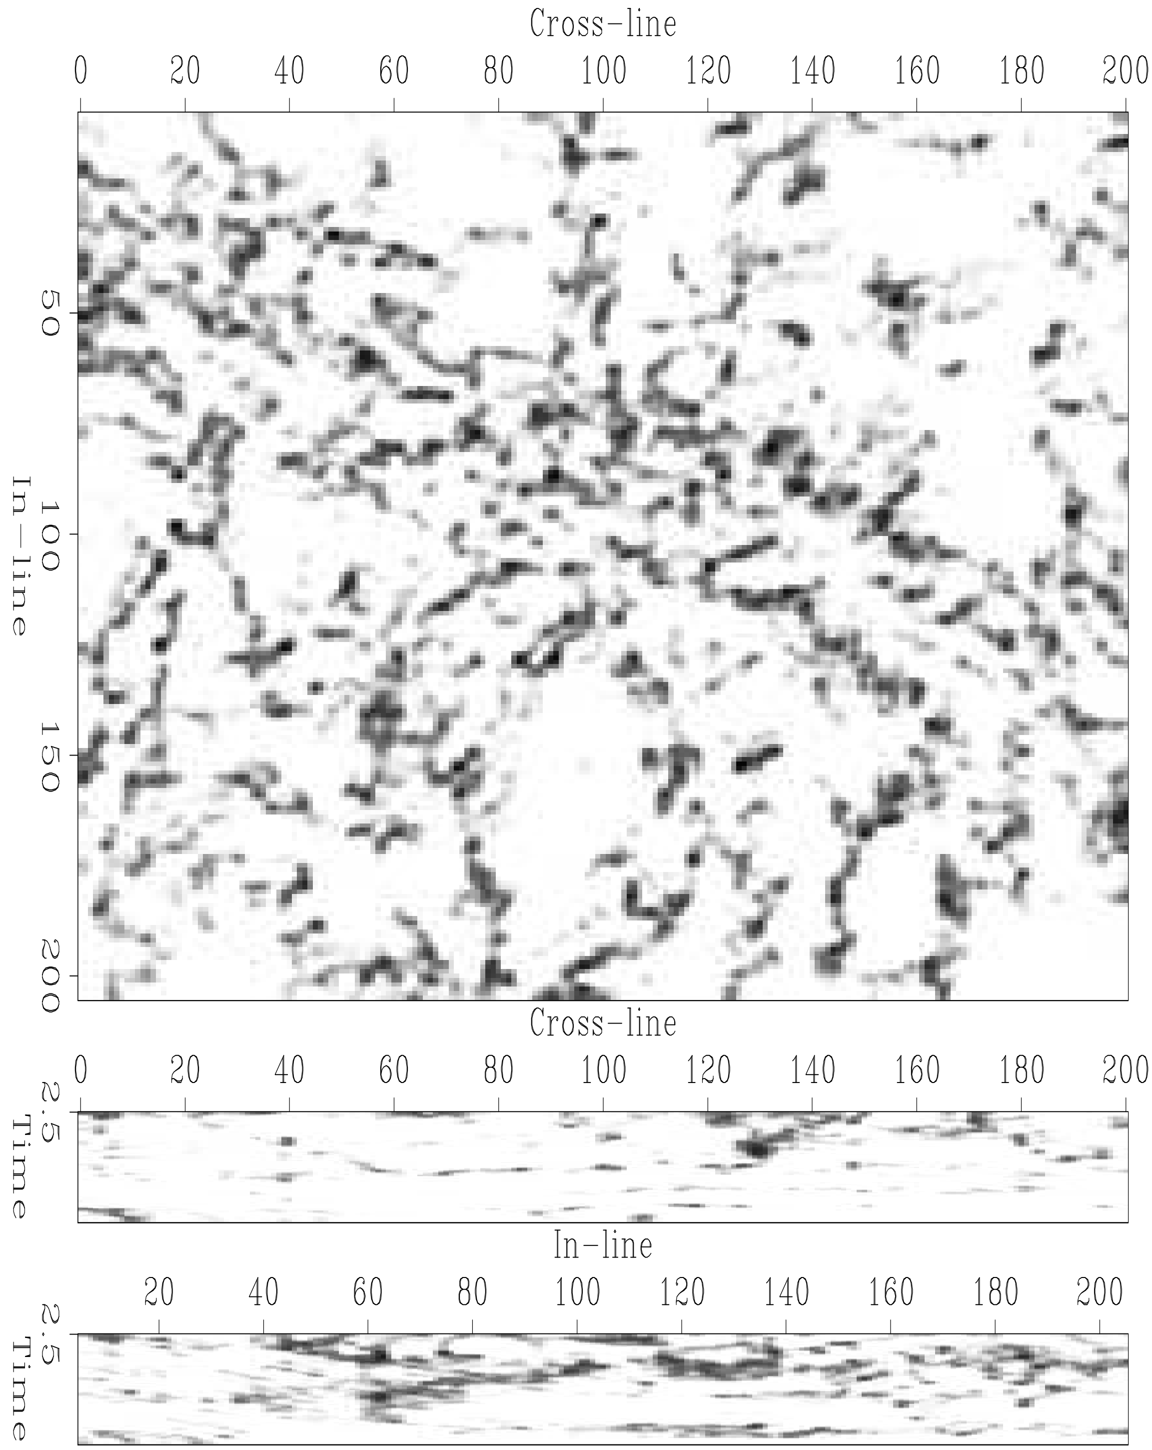

Figure 24: Horizontal correlator applied to horst image. In the time slice, the correlation of neighboring traces successfully removes the sedimentary layers and reveals discontinuities. However, the vertical sections show residuals of the sedimentary layers. nseaFoltTotCorr [CR]

Horizontal correlation, as a semi-successful discontinuity attribute, emphasizes the importance of the within a time slice. The privileged role of the horizontal plane (and its derivatives or correlation) result from the common horizontal deposit of subsurface layers. deposited horizontally. The next section generalizes the horizontal correlation approach to images with sedimentary layers of any dip.

DISCUSSION

As expected, standard edge enhancement techniques, such as the Laplace or the gradient magnitude operator, are unable to discriminate between sedimentary layers or boundaries between layer packages, such as faults.

ACKNOWLEDGMENTS

My colleague Sergey Fomel found the Bay Area Map at a USGS site.

Plane-wave misfit

The misfit of a local plane-wave approximation to a seismic image region – the topic of this chapter – is my most successful discontinuity attribute (Figures 33 and 34). An image location is considered free of a discontinuity if it is locally well-approximated by a plane wave. Discontinuities violate this local plane-wave assumption.

I compute the discontinuity attribute volume in a series of steps. First I partition the image volume into small sub-volumes, *patches*. Next, I estimate the best-fitting plane wave for each patch. The misfit between the plane-wave approximation and the original patch yields a local discontinuity attribute. Finally, I merge all local discontinuity patches into a single discontinuity volume comparable to the input image volume. A later chapter discusses this patching of nonstationary images in greater details.

The three-dimensional plane-wave approximation within a patch estimates the plane-wave slope by minimizing a cross-product expression of partial derivatives. The plane wave's one-dimensional waveform is approximated by stacking along the estimated plane-wave slope. Claerbout (1992) describes an equivalent approach for dip picking in two-dimensional seismic sections. The approach can be extended to a more precise iterative estimation (**?**)

This chapter details the estimation of a best-fitting plane-wave function to a given image patch. The subsequent section discusses various measures of plane-wave fit based on either the residual of the plane-wave dip estimation or the residual of the plane-wave waveform estimation.

PLANE-WAVE ESTIMATION

A three-dimensional plane-wave function is a projection of a one-dimensional waveform $f(s)$:

$$f(\mathbf{p} \cdot \mathbf{x}) = f(p_x x + p_y y + p_z z) \tag{8}$$

where $\mathbf{p}$ is normal to the constant planes. Additionally, $\mathbf{p}$ is non-zero ($||\mathbf{p}|| \neq 0$). Such a projection spreads the value $f(s)$ along the plane perpendicular to the normal direction $\mathbf{p}$. For any given point $\mathbf{x}$, the projection $s = \mathbf{p} \cdot \mathbf{x}$ yields the argument of the function f based on the geometric interpretation of the vector product.

Given an input image $g(\mathbf{x})$, plane-wave estimation finds an approximation of normal vector $\mathbf{p}$ and a waveform estimation $f(s)$. The naive formulation of the plane-wave estimation problem

$$\min_{\mathbf{p}, f(s)} ||g(\mathbf{x}) - f(\mathbf{p} \cdot \mathbf{x})|| \text{ for all } \mathbf{x} \tag{9}$$

is nonlinear and highly non-convex (**?**). Alternatively, the problem can be reformulated into two well-behaved estimations. The estimation

$$\min_{\mathbf{p}} ||\mathbf{p} \times \nabla g(\mathbf{x})|| \text{ for all } \mathbf{x} \tag{10}$$

finds the best-fitting vector $\mathbf{p}$ that is normal to g's gradient. The subsequent estimation

$$\min_{f(s)} ||g(\mathbf{x}) - f(\mathbf{p} \cdot \mathbf{x})|| \text{ for all } \mathbf{x} \tag{11}$$

uses the normal $\mathbf{p}$ estimated by the first minimization to find the waveform $f(s)$.

The formulation of the cross-product minimization stems from the fact that an image volume $g(\mathbf{x})$ is a plane-wave volume with normal $\mathbf{p}$ ($\mathbf{p} \neq \mathbf{0}$), if and only if at all $\mathbf{x}$

$$\mathbf{p} \times \nabla g = \mathbf{0}. \tag{12}$$

If g is indeed a plane-wave function with a normal proportional to $\mathbf{p}$, then

$$g(\mathbf{x}) = f(\alpha \mathbf{p} \cdot \mathbf{x})$$

and, consequently,

$$\mathbf{p} \times \nabla g(\mathbf{x}) = \mathbf{p} \times \nabla f(\alpha\, \mathbf{p} \cdot \mathbf{x}) = \alpha f'(\alpha\, \mathbf{p} \cdot \mathbf{x})\, \mathbf{p} \times \mathbf{p} = \mathbf{0}$$

Similarly, any image $g(\mathbf{x})$ that satisfies equation (12) can be shown to be a plane-wave volume as defined in equation (8).

The constraint of a nonzero normal vector $\mathbf{p} \neq 0$ can be implemented in various ways. The constraint $|\mathbf{p}| = 1$ leads to a symmetric formulation in all spatial variables x, y, and z. However, for mathematical convenience, I choose the constraint $p_z = 1$, which excludes vertical plane waves and leads to an asymmetry between the spatial variables. In practice, this particular constrain does not limit the processing of ordinary seismic images since image features are rarely vertical.

The minimization of equation (12) can be solved analytically. Let $\mathbf{g}_x$ be the vector of partial x-derivatives of discretized g. Each element represents the derivative at a given location $\mathbf{x}$. The vectors $\mathbf{g}_y$ and $\mathbf{g}_z$ are the corresponding partial y- and z-derivatives. Given the constraint $p_z = 1$, equation (12) can be expressed in matrix form as

$$\begin{pmatrix} 0 & \mathbf{g}_z \\ \mathbf{g}_z & 0 \\ \mathbf{g}_y & -\mathbf{g}_x \end{pmatrix} \begin{pmatrix} p_x \\ p_y \end{pmatrix} = \begin{pmatrix} \mathbf{g}_y \\ \mathbf{g}_x \\ 0 \end{pmatrix}.$$

The minimization's corresponding normal equation is

$$\begin{pmatrix} \mathbf{g}_y^2 + \mathbf{g}_z^2 & -\mathbf{g}_x \cdot \mathbf{g}_y \\ -\mathbf{g}_x \cdot \mathbf{g}_y & \mathbf{g}_x^2 + \mathbf{g}_z^2 \end{pmatrix} \begin{pmatrix} p_x \\ p_y \end{pmatrix} = \begin{pmatrix} \mathbf{g}_x \cdot \mathbf{g}_z \\ \mathbf{g}_y \cdot \mathbf{g}_z \end{pmatrix}. \tag{13}$$

The variables

$$q_x = \frac{\mathbf{g}_x \cdot \mathbf{g}_z}{\mathbf{g}_z \cdot \mathbf{g}_z} \text{ and } q_y = \frac{\mathbf{g}_y \cdot \mathbf{g}_z}{\mathbf{g}_z \cdot \mathbf{g}_z} \tag{14}$$

estimate the gradient slope in the x and y direction of g. Substitution of the gradient expressions (14) into the normal equation (13) yields

$$\begin{pmatrix} q_y^2 + 1 & -q_x q_y \\ -q_x q_y & q_x^2 + 1 \end{pmatrix} \begin{pmatrix} p_x \\ p_y \end{pmatrix} = \begin{pmatrix} q_x \\ q_y \end{pmatrix}.$$

Finally, the equation system can be solved for the x- and y-component of the plane wave's normal vector:

$$\begin{pmatrix} p_x \\ p_y \end{pmatrix} = \frac{1}{(q_x^2+1)(q_y^2+1) - (q_x q_y)^2} \begin{pmatrix} q_x^2 + 1 & q_x q_y \\ q_x q_y & q_y^2 + 1 \end{pmatrix} \begin{pmatrix} q_x \\ q_y \end{pmatrix}.$$

Claerbout (1992) shows how the gradient slopes q_x and q_y can be computed from the input image using finite-difference approximations of the partial derivative expressions $\partial g/\partial x$, $\partial g/\partial y$, and $\partial g/\partial z$. The solution fully specifies the normal vector $\mathbf{p}$ since the vector's z component p_z is constrained to 1.

If we set the variables in the y dimension – q_y and p_y – to zero, the solution simplifies to

$$p_x = q_x = \frac{\mathbf{g}_x \cdot \mathbf{g}_z}{\mathbf{g}_z \cdot \mathbf{g}_z},$$

which is identical to Claerbout's (1992) two-dimensional dip picking scheme. My original cross-product formulation (12) expresses the natural symmetry of the plane-wave detection problem that Claerbout's Lomoplan formulation lacks. However, my constraint for the nonzero normal vector ($p_z = 1$) introduces the same asymmetry later.

Stacking of the image volume $g(\mathbf{x})$ along planes orthogonal to the estimated normal vector $\mathbf{p}$ yields a first estimate of the unknown plane-wave waveform $f(s)$. Symes (**?**) formulates a similar, but more sophisticated, iterative estimation scheme of the normal vector $\mathbf{p}$ and the waveform $f(s)$.

DIP MISFIT BY CROSSPRODUCT RESIDUAL

The residual of the least-squares estimation (12)

$$\mathbf{r_p}(\mathbf{x}) = \mathbf{p} \times \nabla g(\mathbf{x})$$

indicates how well the seismic volume is approximated by a plane wave. The computation does not explicitly involve the plane-wave waveform $f(s)$, but is the cross-product of the the local gradient ∇g and the patch's best-fitting plane-wave gradient $\mathbf{p}$. For a three-dimensional input volume, the residual of the cross-product operator expression yields a three-dimensional vector at each output location of the three-dimensional image volume.

Results

Figure 25 shows the cross-product output for the synthetic test case of Figure 4. Each component of the local residual vectors is displayed as an individual image volume. In the regions away from the fault, where the image satisfies the plane-wave assumption, the residual is zero. At the fault, the output image is non-zero and the fault is broadly outlined. Within a patch that straddles the fault, the plane-wave dip estimation encounters two contradictory plane-wave events. The least-squares solution averages the two dips so that the dip estimate does not approximate either of the two dips. Consequently, within a patch that straddles the fault, all pixels show a significant residual. The patch-wide residual reduces the resolution of the discontinuity image to the size of the patches and does not imply a meaningful interpretation of the discontinuity amplitude.

Because a three-dimensional vector field is difficult to interpret, the schemes of the two next sections summarize the information of the three volumes into a single discontinuity image.

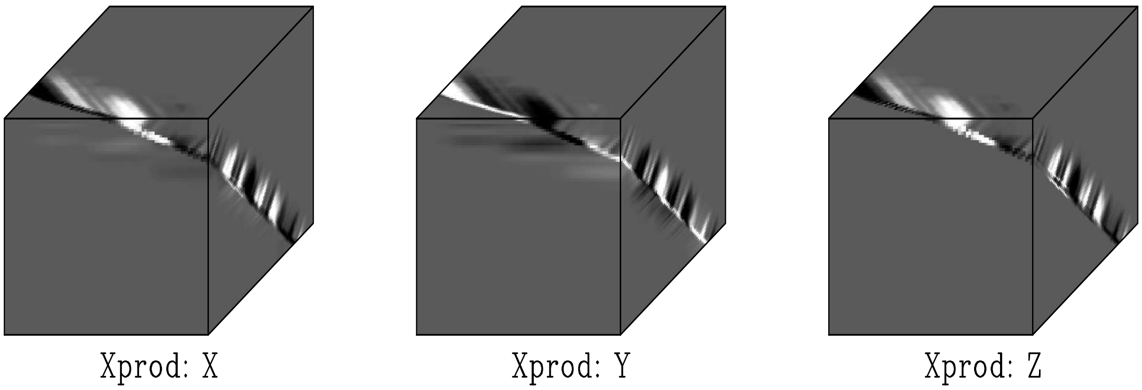

Figure 25: Cross-product operator applied to synthetic image. The cross-product operator applied to the synthetic fault model broadly delineates the fault in each of its three output volumes. Each output volume is a component of the cross-product vector. However, the three-dimensional output vector at each pixel of the three-dimensional image is not suited to interpretation. zeroFoltDXG [ER]

NORM OF RESIDUAL

The magnitude of the three-dimensional residual vector $\mathbf{r_p}$ at each patch pixel location $\mathbf{x}$,

$$r_L(\mathbf{x}) = \sqrt{\sum_{\mathbf{p}} \mathbf{r_p}(\mathbf{x}) \cdot \mathbf{r_p}(\mathbf{x})},$$

collapses the vector field $\mathbf{r_p}$ to an image of the same size and shape as the input. The scalar r_L is zero if and only if the local $\mathbf{r_p}$ is the null vector, which in turn implies that the local gradient $\nabla g(\mathbf{x})$ is parallel to the patch's best-fitting plane-wave normal $\mathbf{p}$.

Results

The magnitude of the residuals of the synthetic image example (Figure 26) resembles its constituent components (Figure 25). Patches that contain pure plane waves are zeroed. Patches that straddle the fault and contain two plane waves of distinct dip show significant residual amplitudes. The limit of resolution of the operator is the patch and, again, the fault is broadly outlined but not pinpointed. The discontinuity amplitude does not have a simple interpretation, except that it indicates plane-wave elements that are not well approximated by the somewhat arbitrary best-fitting local plane wave. Interestingly, the generally lower amplitudes of the horizontal plane-wave region indicate that the the best-fitting local plane waves tend to approximate the horizontal events better than the dipping ones. The side of the local cubic patch is 12 pixels long; the side of the total input cube is 80 pixels long.

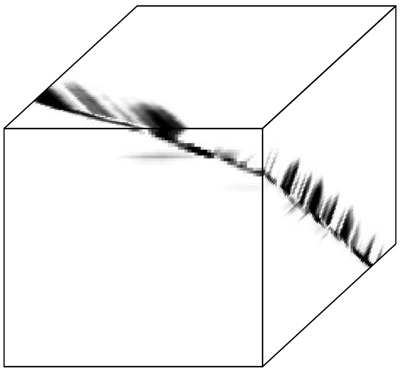

Xprod: Pixel magnitude

Figure 26: Pixel magnitude of cross-product residual of the synthetic image. The pixel magnitude of the cross-product residual of Figure 25 roughly delineates the fault of the input image. Again, instead of pinpointing the fault, the operator highlights any local patch that violates the plane-wave assumption. zeroFoltDXGLN [ER]

Applied to the Gulf salt dome image (Figure 27), the magnitude of the crossproduct operator generates a noisy discontinuity map that is void of sedimentary layers but only hints at the image's faults. The attribute suppresses the sedimentary layers successfully, even the steeply dipping ones in the north-west corner. The radiating faults (e.g., the one marked Rin the original image 8) are depicted as linear features of high discontinuity amplitude. However, the image's general random noise obscures its faults. The internal area of the salt body does not generate high amplitudes. The pentagonoid region within the salt, the diagonally truncating fault, and the salt-sediment boundary are broadly outlined by a fuzzy zone of high discontinuity.

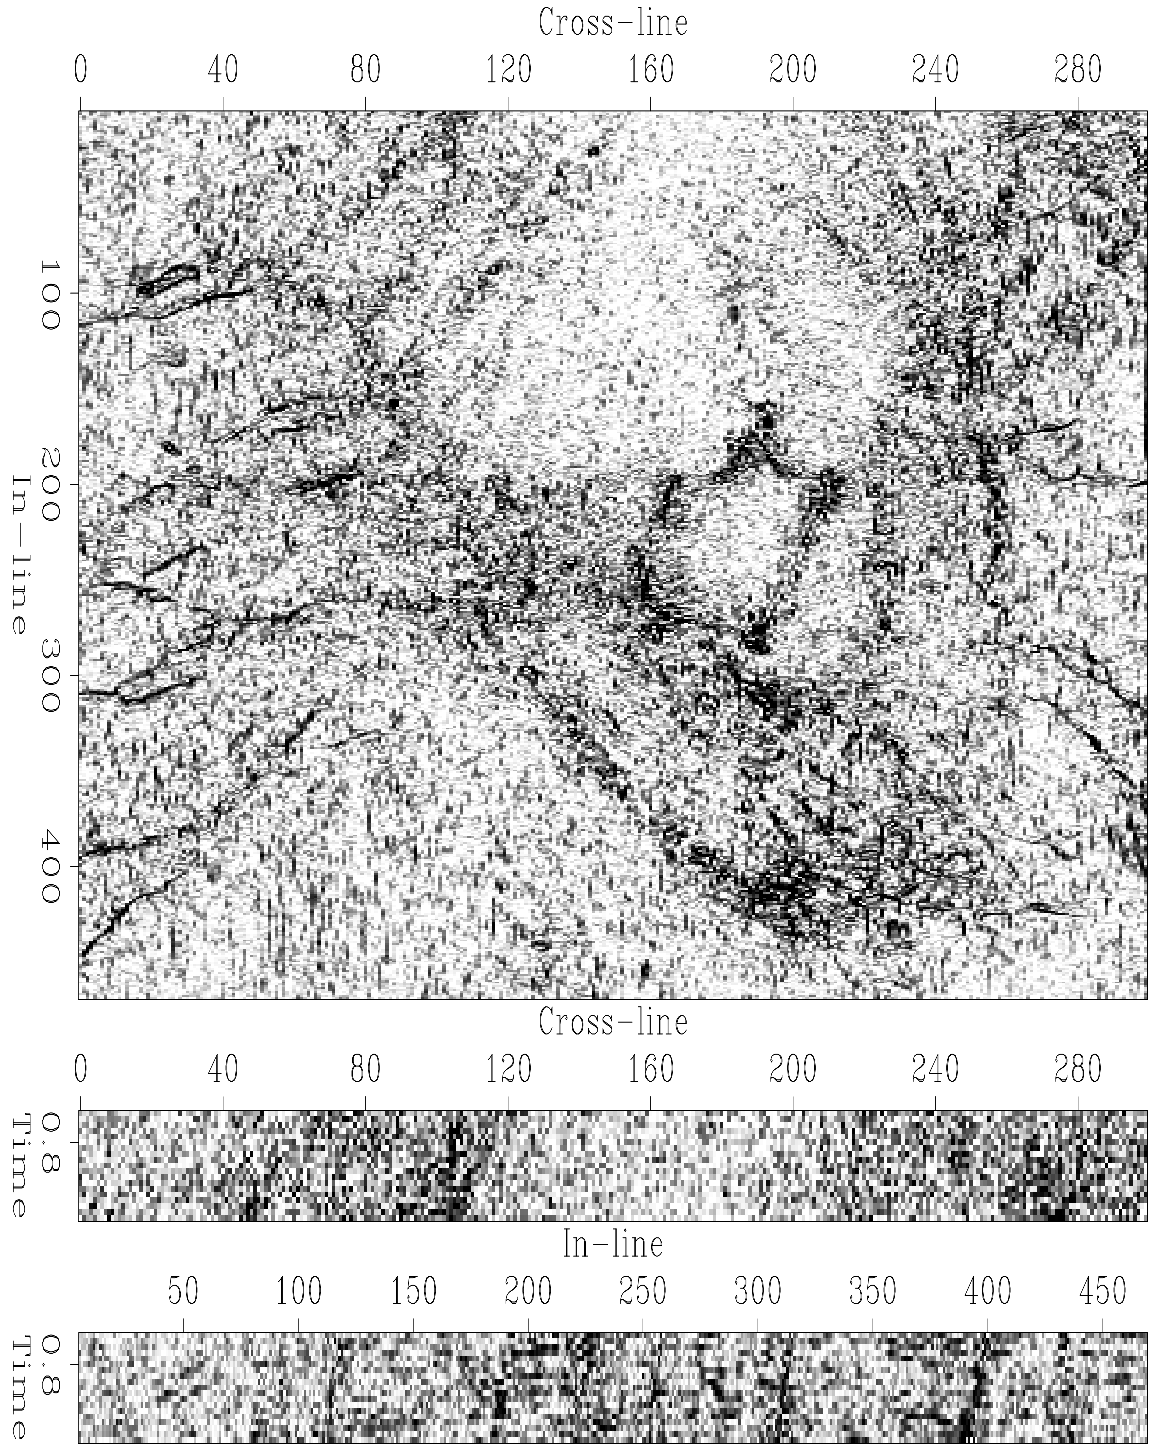

Figure 27: Pixel magnitude of cross-product residual of the Gulf salt image. The noisy discontinuity map successfully suppresses the original wavefield character and fuzzily indicates the image's faults. gulfFoltTotDXGLN [CR]

The interpretability of the image features suffers from its high noise-to-signal ratio that spoils its contrast and sharpness. The resolution of the radial faults is again the size of the patches. Moderate temporal smoothing might slightly improve the image without loss of resolution.

In the case of the North Sea horst image (not shown), the randomness of the time slice defeats interpretation, even though a few linear features are discernible. Disappointingly, the vertical section of the North Sea horst image depict traces of sedimentary layers. In summary, the magnitude of the cross-product residual failed to compute an interpretable discontinuity map.

BACKPROJECTION OF RESIDUAL

The compounding of the cross-product operator (12) and its adjoint yields an alternative discontinuity measure, the back-projected residual:

$$r_A(\mathbf{x}) = (\mathbf{p} \times \nabla) \cdot \mathbf{r_p}(\mathbf{x}) = (\mathbf{p} \times \nabla) \cdot (\mathbf{p} \times \nabla) g(\mathbf{x})$$

The output and input image agree in size and shape. In contrast to all-positive norm expressions, the back-projected residual can be negative. Unfortunately, I do not have an insightful interpretation of the back-projected residual.

Results

The backprojected residual of the synthetic image example (Figure 28) resembles the individual components of the crossproduct operator. The fault is outlined but directly bordered by large residual amplitudes of the neighboring plane wave regions. Consequently, the resolution is about the size of the patch. Otherwise, the plane wave patches away from the discontinuity are zeroed.

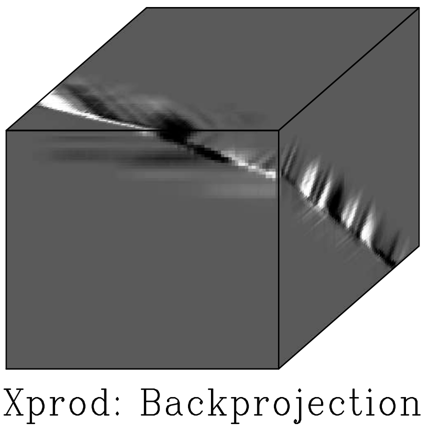

Figure 28: Forward and adjoint of cross-product operator applied to the synthetic image. The back-projected residual of the synthetic test image suppresses the pure plane-wave patches of the image and broadly delineates the input's fault location. zeroFoltDXGAA [ER]

In the case of the Gulf salt dome image (Figure 29), the back-projecting crossproduct operator removes the image's sedimentary layers successfully, but the delineated faults are slightly corrupted by the image's considerable noise. Some of the noise forms short north-south streaks – best visible in the triangular area in the south-west corner – and might be due to the image's acquisition footprint. The border of the central low-amplitude region is outlined by a broad and fuzzy zone of higher average discontinuity amplitudes. The salt region shows some internal features, such as the linear north-south streaks in its center. The salt's pentagonoid region is visible but its western border is strangely distorted. The salt image's vertical north-south section shows a disappointing amount of remaining sedimentary layers. In contrast, the east-west section shows only a few weak hints of the former layering. Naturally, the near-vertical faults are more visible among the coherent layers of the north-south section than the rather incoherent east-west section.

The back-projected crossproduct operator applied to the North Sea horst image generates an extremely noisy discontinuity map. A few linear features can be seen at a grazing angle. The features include the fault marked `F` in the original image 12. The sparsity of features contrasts positively with the complexity of the same time slice when processed by competing discontinuity operators. However, the fault surfaces are not isolated in the image noise, and the vertical sections show considerable layering.

In summary, the back-projection operator highlights many interesting details that other discontinuity attributes miss. However, its overall low signal-to-noise ratio does prevent interpreters from quickly and conveniently assessing the geological information.

PLANE-WAVE MISFIT BY CORRELATION

The direct comparison between a local image patch and its best-fitting plane wave offers another alternative discontinuity attribute. How do I estimate the best-fitting plane wave $\hat{f}(\mathbf{p} \cdot \mathbf{x})$ of an image region $f(\mathbf{x})$? In the previous section, I showed how to estimate the best-fitting plane-wave normal $\mathbf{p}$ for a given image region. Given its normal estimate $\mathbf{p}$, stacking the image f along the hyper-planes orthogonal to $\mathbf{p}$ yields an estimate of the one-dimensional waveform $\hat{f}(s)$. The adjoint of the stacking procedure yields an estimate of the plane-wave volume $\hat{f}(\mathbf{p} \cdot \mathbf{x})$. The stacking operator and its adjoint use weights to compensate for the variable number of contributions for each waveform element.

The plane wave $\hat{f}(\mathbf{p} \cdot \mathbf{x})$ is not best-fitting the original $f(\mathbf{p} \cdot \mathbf{x})$ in a least-square sense of the expression (9). Instead, the normal $\mathbf{p}$ is a least-squares approximation of the true normal. The waveform is an estimate that depends on the normal $\mathbf{p}$ (rather than being estimated simultaneously). If the original image is a plane wave, my process, of course, yields the correct estimate.

The reformulation of the problem greatly simplifies the plane-wave estimation. I reformulate a highly nonlinear, non-convex cost function into a linear estimation of the plane-wave normal $\mathbf{p}$ and the linear estimation of the waveform f.

Many geophysical problems are conveniently expressed as optimization problems. Unfortunately, the computational costs of these optimization problems (often expressed by the number of iterations necessary to converge) deliver these optimal formulations often

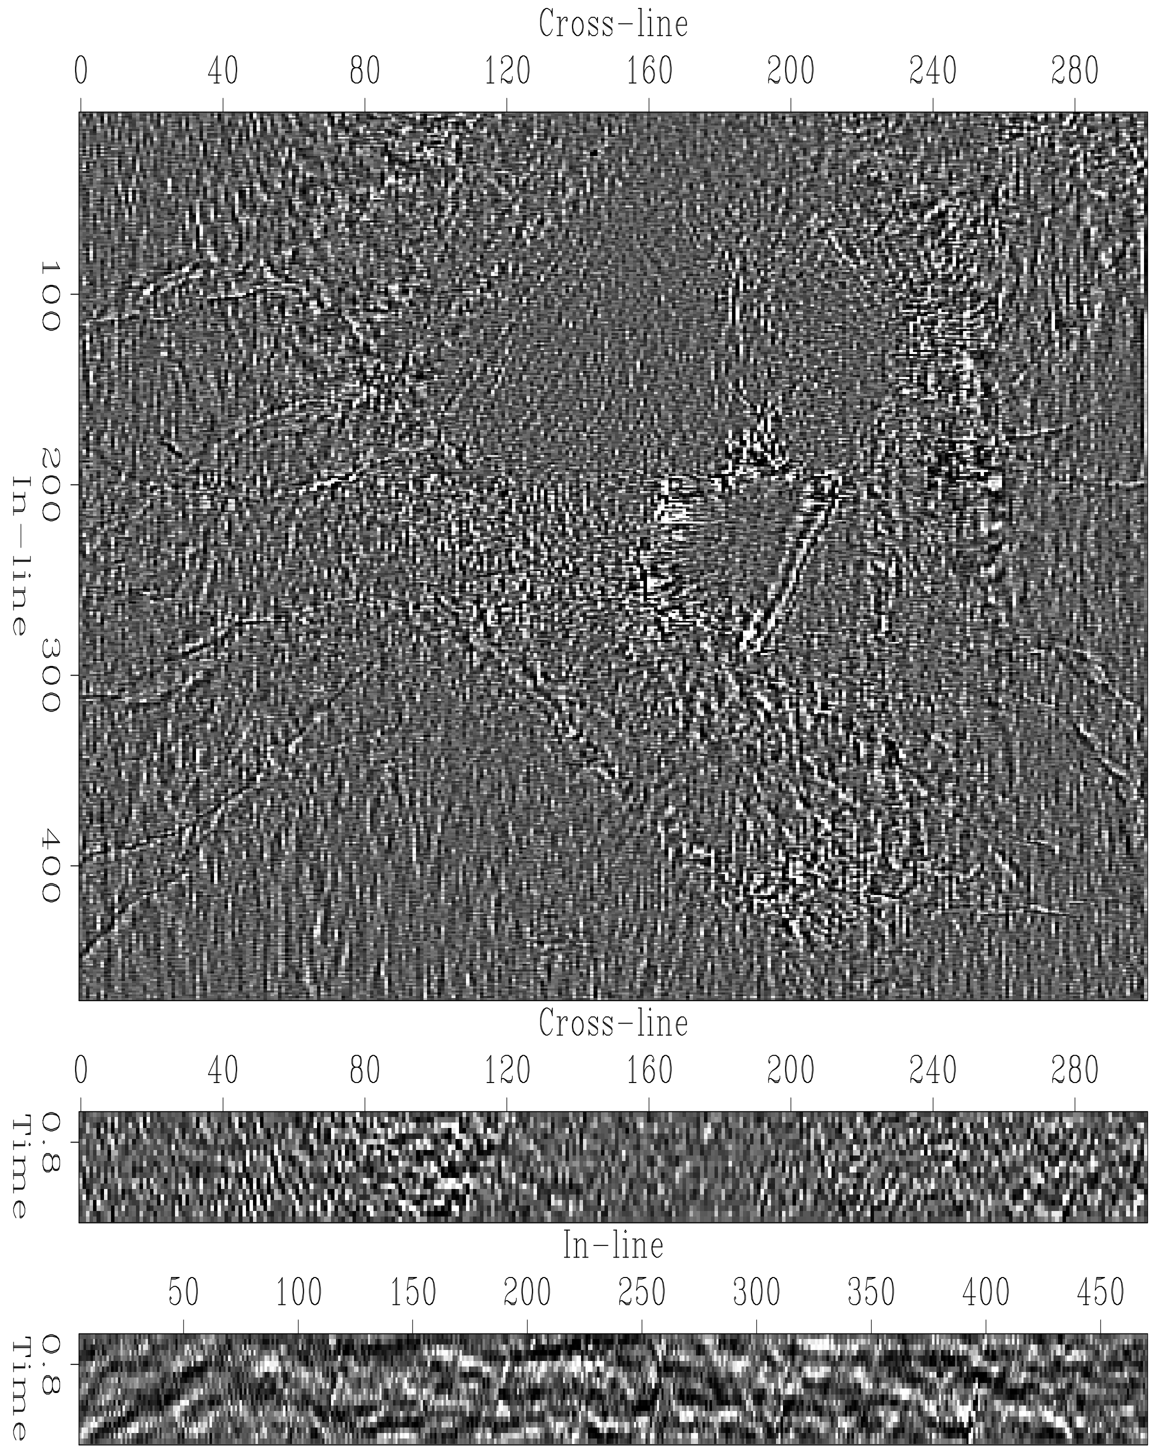

Figure 29: Back projection of cross-product residual of the salt image. The discontinuity map successfully suppresses the layers in the time slice and delineates the radiating faults. Overall, the image's noise makes it unattractive for interpretation. gulfFoltTotDXGAA [CR]

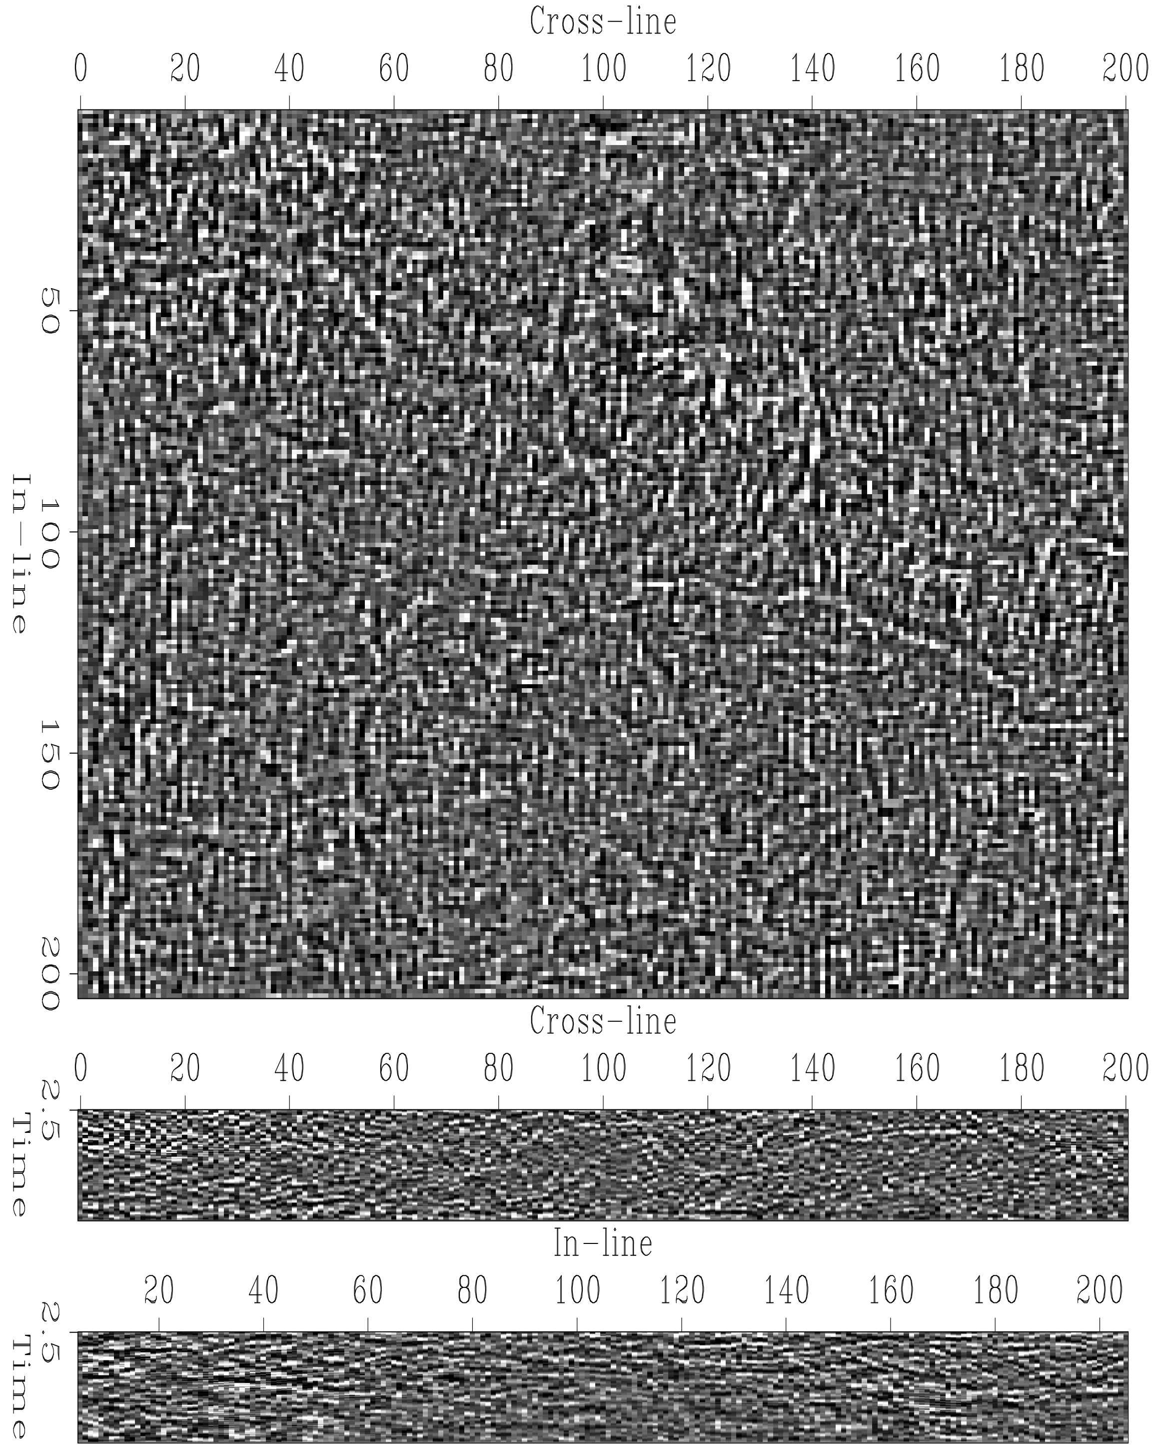

Figure 30: Back projection of cross-product residual of the North Sea horst image. The image is dominated by noise. At a grazing angle, the time slice depicts a few gently bending linear features (faults?). nseaFoltTotDXGAA [CR]

unrealistic for routine computations. I believe, a central challenge for geophysicists today is the mathematical reformulation of geophysical optimization problems into computationally manageable procedures. The crossproduct formulation of the plane-wave estimation problem is such a reformulation.

Figure 31 displays the composite image of the local best-fitting plane waves for the synthetic test case. I subdivide the original synthetic data of Figure 4 into local patches without overlap, estimate in each patch the best-fitting local plane wave, and merge the local plane waves to a composite image. The plane-wave components in the dipping and horizontal half-space are correctly estimated. Along the boundary of the two half-spaces, however, the synthetic data deviates from the single plane-wave assumption. The least-squares estimate of the dip averages the correct values of either side. The combination of stacking and its adjoint nevertheless fills each patch with a plane wave volume, which, of course, does not resemble the original patch contents of two plane waves.

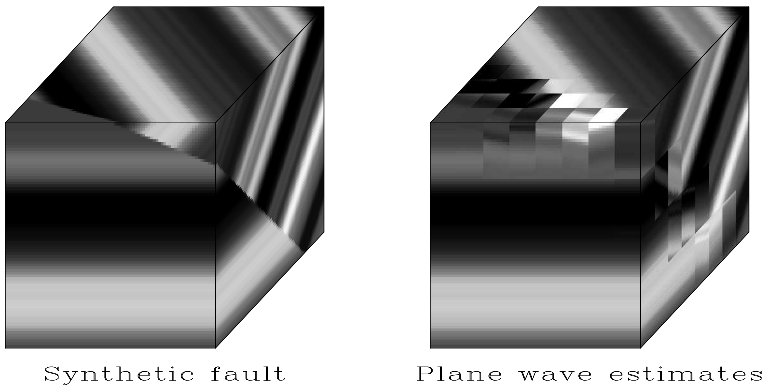

Figure 31: Locally best-fitting plane-wave volumes in the synthetic image 4. Away from the fault surface, each patch's plane wave is correctly estimated. For patches along the fault surface, the single plane-wave assumption is violated and the dip and local wavefield estimates are incorrect. zeroFoltPP [ER]

To compare the best-fitting plane wave and the original local image region I correlate the patches' corresponding individual traces (the patches' individual time intervals). The discontinuity attribute divides the image into patches, estimates the best-fitting plane wave of each patch, computes a correlation coefficient for each trace of a patch, sets all samples of a trace to the correlation coefficient, merges such patches – now indicating correlation – into a single image volume.

The correlation of two traces is

$$c(f, \hat{f}) = \frac{\sum_i f_i \hat{f}_i}{\sqrt{\sum_i f_i^2 \sum_i \hat{f}_i^2}}$$

where i loops over the samples of the one-dimensional trace f and $\hat{f}$. The coefficient is bounded by $[-1, 1]$ and yields 1 only if the two traces are identical.

In general, correlation is a robust estimate of similarity as it averages over an ensemble of samples. The discontinuity attribute could loop over all kind of subsets of an individual patch to compute correlation. The entire patch could be thought as a supertrace, which would be consistent with the plane-wave estimation of a single patch. I choose a single coefficient per trace since it averages over time samples and, thereby, looses vertical resolution, but preserves lateral resolution. Such a tradeoff seems to be well-suited for seismic data, which is usually oversampled along the time axis, but sparsely sampled in space, and which often depicts horizontal strata. Nevertheless, the overall lateral resolution of the attribute is limited by the attributes' ability to estimate the normal and waveform at discontinuities (as described above).

As with the display of the horizontal correlation attribute, I simplify the attribute maps by linearly transforming the correlation values to non-negative numbers. The optimal correlation of 1 is displayed as 0 amplitude, the better correlated the image, the smaller the attribute value. In comparison, the correlation (6) in this section estimates the similarity between two one-dimensional arrays. The correlation (7) used in the horizontal correlation coefficient, on the other hand, estimates the similarity *among a set* of one-dimensional arrays.

Results

Overall, the correlation of the best-fitting plane wave is the my most successful discontinuity attribute. The synthetic test case shown in Figure 32 overall delineates the fault. At the cube's side views, the elongated events that constitute the fault indicate the attribute's loss of vertical resolution. In contrast, the fault is sharply defined in the top view. Unfortunately, the increased lateral resolution causes the fault line to be interrupted where adjacent layers lack contrast.

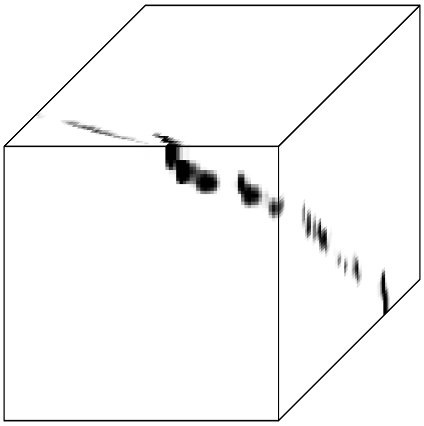

Figure 32: Correlation of local plane-wave estimate and original image. The fault is resolved sharply laterally, but is blurred vertically. The sharp lateral resolution leads to gaps in the fault line where the original adjacent layers lack contrast. zeroFoltPPC [ER]

The first seismic test case in Figure 33 is a success. The wavefield character of the image due to the sedimentary layers vanished. The major faults are delineated, particular the large fault in the south-west corner. The faults are sharply resolved, and contrast well with the otherwise eventless background.

Unfortunately, the image shows a few shortcomings. In some locations, the image shows a distinct ringing that seems to be due to insufficiently suppressed, steeply dipping layers. The salt body and its internal chaotic reflections obscure the center of the image. The pentagonoid region is not delineated, nor are the major truncating faults within the salt body. In places, the image shows some distinct north-south streaks, which I believe are due to the acquisition footprint that is also visible in the original image.

The loss of resolution and lack of faults in the vertical sections of Figure 33 is due to the correlation over time. The distinct border between the upper and lower data part indicates that the patch size should possibly be reduced.

Overall, the quality of the discontinuity image seems comparable to the quality of commercial discontinuity maps (Bahorich and Farmer, 1995; **?**). My salt dome image is a subimage of one of Bahorich's examples and the attributes can be compared directly.

The discontinuity image of Figure 34 suppresses the wavefield character of the original image, which was not excessive in the first place. The discontinuity map reveals a complex pattern of numerous faults and discontinuities. The discontinuities are well resolved. The discontinuity map is slightly contaminated with noise, which reduces the sharpness and contrast of the individual events and which is most-visible in the image's center. The noise might be due to events that only partially fall within the correlation's temporal averaging. The vertical section indicate that a few strong sedimentary layers at the bottom were not suppressed well.

Are the image's discontinuities truthfully delineated? The faults that I can identify in the original image all correspond to events of the discontinuity map. Some of the events – e.g., the fault marked `F` in the original image 12 – deviate from the rather continuous and smooth event that I detect in the original seismic image. Other equally strong events of the discontinuity image seem to relate to what amounts to small, hardly visible discontinuities in the original image. If my observations are correct, one may argue that the discontinuity attribute fails to distinguish between major and minor discontinuities. On the other hand, the method deserves credit for detecting such minor discontinuities.

DISCUSSION

The trace correlation between an image region and its best-fitting plane-wave yields an easily interpreted seismic section. The resolution of the discontinuity map is about the size of the local patch that is used to estimate the best-fitting plane wave. The quality of the discontinuity map is comparable to similar to Amoco's original coherency attribute. In comparison to the correlation, alternative misfit measures of the locally best-fitting plane wave yielded inferior discontinuity maps.

Plane-wave estimation

The plane-wave estimation by my cross-product expression is both inexpensive and elegant. I believe the formulation enables us to routinely use plane-wave estimation in a wide variety of processing applications. I am particularly interested in the fitting of the dip of seismic reflections and of well log readings. Another application is the migration of sparse seismic data using reflection dip.

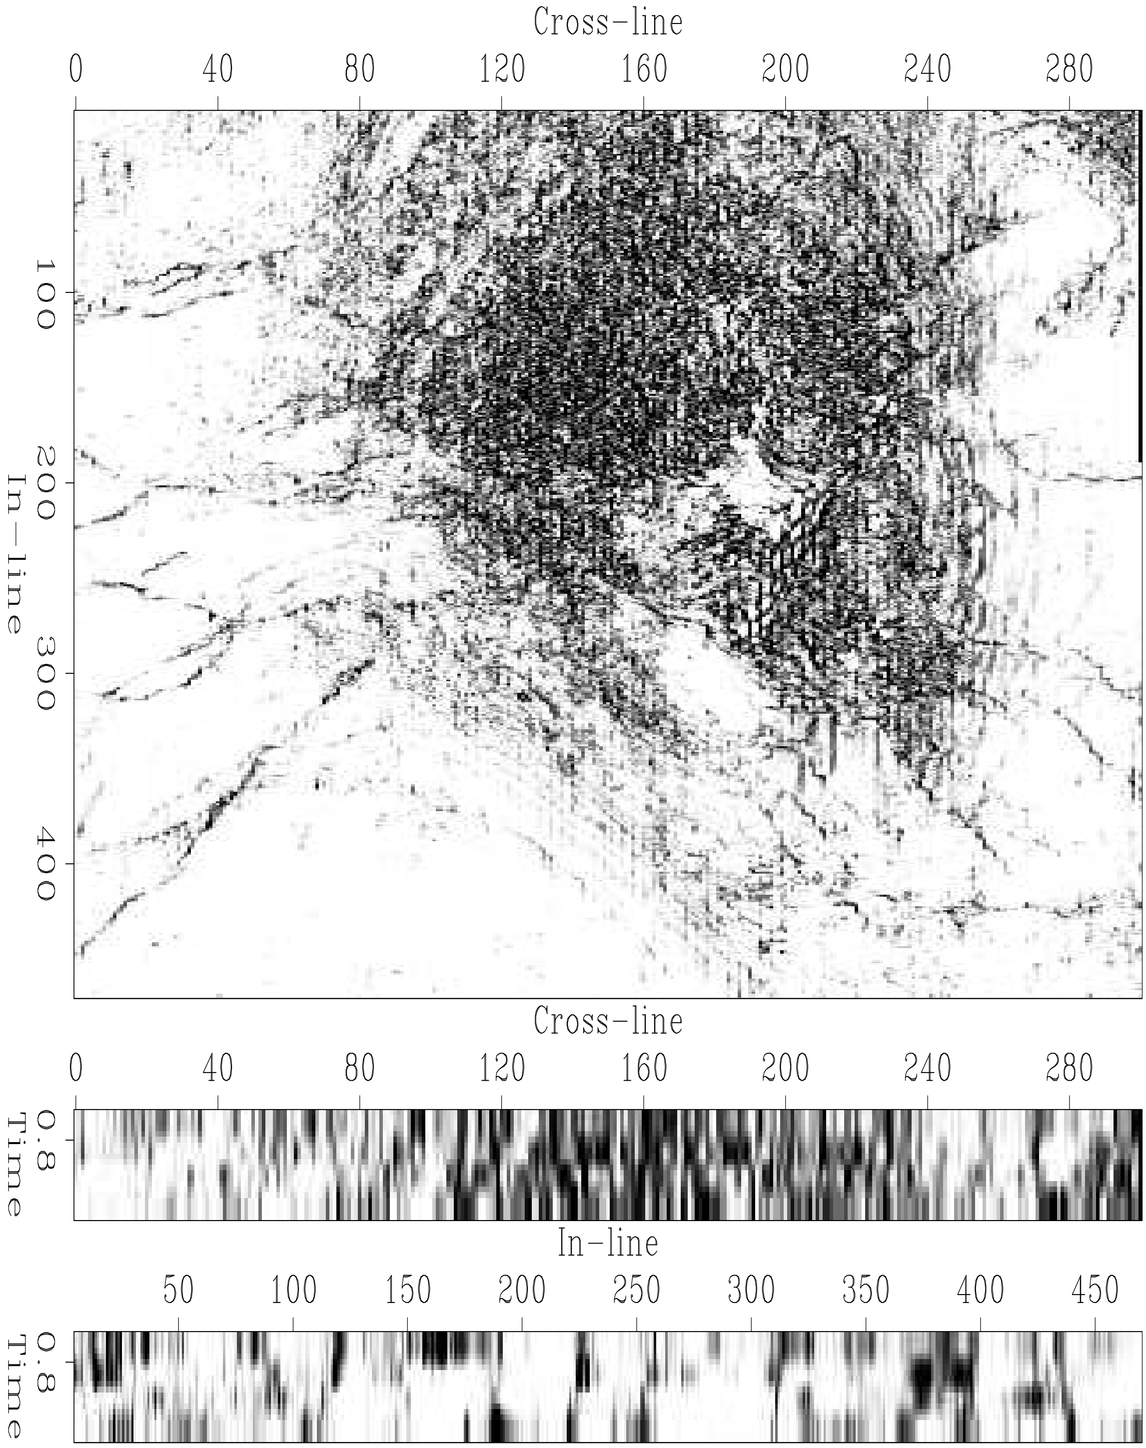

Figure 33: Plane-wave correlation applied to the salt image. The image's time slice successfully delineates the subsurface faults and is easily interpreted. $\boxed{\text{gulfFoltTotPPC}}$ [CR]

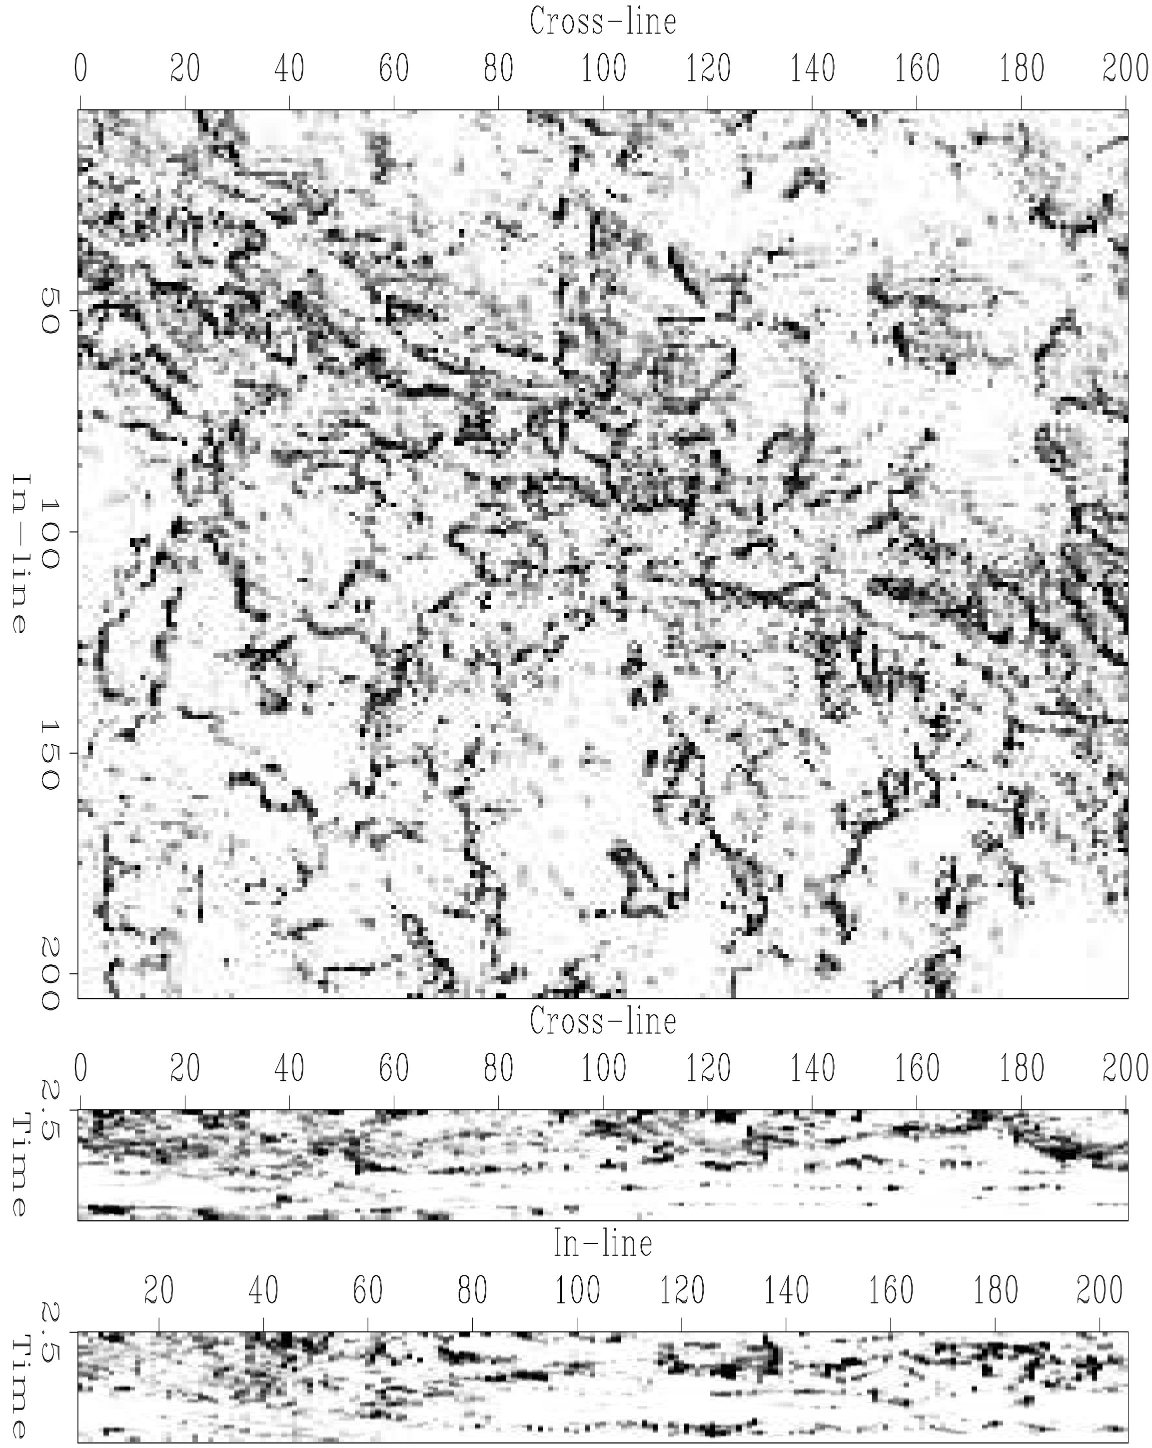

Figure 34: Plane-wave correlation applied to the horst image. The discontinuities of the image are delineated and the image offers a distinctively different and informative view of the subsurface. nseaFoltTotPPC [CR]

ACKNOWLEDGMENTS

Jon Claerbout previously developed a two-dimensional finite-difference dip estimation scheme that turned out to be related to the three-dimensional one I discovered. Bee Bednar urged me to combine *somehow* two of Jon's operators to a three-dimensional discontinuity attribute and, thereby, prodded me into exploring the plane-wave estimation in greater detail than I first intended. Bill Symes pointed out the relationship between the cross-product operator and his Differential Semblance Optimization work.

Prediction error

Prediction error is a discontinuity attribute that removes the predictable image components and reveals the unpredictable. To use prediction error as a discontinuity attribute – the original goal and starting point of my project – one has to devise a prediction-error computation that predicts and removes the plane-wave volumes of sedimentary layers but that is incapable of predicting the discontinuities.

Prediction-error filters (Claerbout, 1992; Jain, 1989) remove the image component that is predictable as a linear combination of its neighboring values. Given a two-dimensional image $u(m,n)$ and the linear prediction estimate

$$\hat{u}(m,n) = \sum_{k \in S_x} \sum_{l \in S_x} a(k,l)\, u(m-k, n-l),$$

the prediction error is simply $\epsilon(m,n) = u(m,n) - \hat{u}(m,n)$.

The region S_x defines which neighboring values contribute to the linear prediction. Causal predictions, that involve regions of the shape shown in Figure 35, lead to white noise driven output images. The region is called *causal* since given a hypothetical scan from top to bottom and left to right all points of S_x lie to one-side of the predicted value $u(m,n)$.

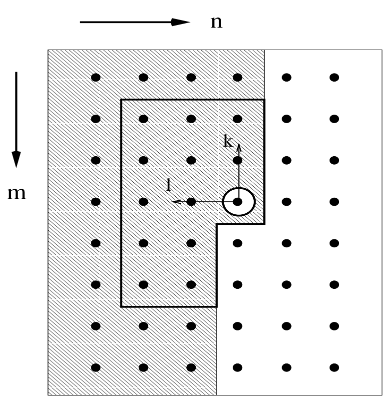

Figure 35: Prediction-error domain. The causal domain ensures that the output of the prediction-error filter tends to be white noise. pefDomain [NR]

To compute the prediction error of a given stationary image, we first find the prediction coefficients $a(k,l)$ that minimize the prediction error $\epsilon(m,n)$ for all pixels of the input image. Once the prediction coefficients are known, the convolution

$$\epsilon(m,n) = u(m,n) - \sum_{k \in S_x} \sum_{l \in S_x} a(k,l)\, u(m-k, n-l).$$

computes the prediction error. In particular, a prediction-error filter potentially zeroes a random plane-wave, or a superposition of random plane waves, or a superposition of random constant-amplitude lines. I represent the prediction-error computation as

$$\epsilon(\mathbf{x}) = A\, g(\mathbf{x}),$$

where g is the input image, A is the prediction-error operator, and ϵ the residual prediction error.

To compute the prediction error of a nonstationary image, I, as usual, divide the image into stationary patches, compute the prediction error for each patch, and merge the patches containing the local prediction error into a single quilt. All result images of this section are smoothed along the vertical axis to suppress the prediction error's tendency to enhance high-frequency noise of the original unfiltered image.

THREE-DIMENSIONAL PREDICTION-ERROR FILTER

A three-dimensional prediction-error filter applied to a stationary image volume will tend to whiten it unconditionally. Potentially, any coherent component of a stationary image is removed. In practice, the size of the filter and the accuracy of the estimation of its coefficients limit the predictive powers of a filter and yield an image that is not entirely white. In particular, a three-dimensional prediction filter removes a plane-wave image.

Results

Figure 36 shows that a single three-dimensional prediction-error filter applied to the synthetic test image suppresses the plane-wave events in the patches away from the fault surface. For patches that straddle the fault, the prediction error is considerable. In these nonstationary patches, the prediction coefficients are optimized to remove plane waves of two different dip simultaneously. However, since the plane-wave events are not superimposed but adjacent, the filter is unable to predict either event perfectly. The unpredictable nature of neighboring amplitudes along the discontinuity yields particularly high amplitudes.

The horizontal events tend to be better predicted than the dipping events, probably due to round-off errors.

Figure 37 shows the prediction-error attribute of the salt dome image. The image's original wavefield character is completely suppressed. The image indicates the faults, but their exact location is obscured and interrupted by noise. Due to the noise, the image lacks sharpness and contrast. The detected faults are well resolved, but difficult to delineate. In particular, the fault marked `R` in the original image 8 is blurred at its fork. The busy

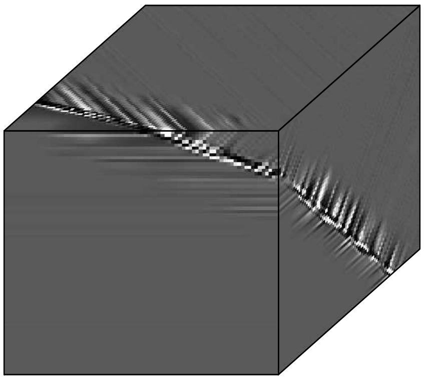

Figure 36: 3-D prediction-error of synthetic image. The three-dimensional prediction-error filter predicts and removes the plane-wave events in patches undisturbed by the discontinuity. The patches that straddle the discontinuity are only partially removed. zeroFolt1Pef3d [ER]

region at the southern tip of the salt body hides the local faults. Within the salt, the attribute barely hints at the pentagonoid region. The salt itself is characterized by incoherent low amplitudes. The salt boundary is a zone of increased noise level. The boundary is diffuse and does not separate salt and adjacent sediments authoritatively. Overall, the attribute is too complex and noisy to benefit geological interpretation.

At a first impression, the prediction error of the North Sea horst image (Figure 38) appears to be white noise. Only careful inspection – especially at a grazing angle – reveals the image's complex discontinuity pattern. However, the faults are not easily delineated.

RESIDUAL OF THREE TWO-DIMENSIONAL PREDICTION-ERROR FILTERS

The three two-dimensional one-column prediction-error filters shown in Figure 39 predict and remove a plane-wave volume (stacked pancake model). Any other image then a plane-wave volume leads to nonzero prediction error in at least one of the three output images. In contrast to a three-dimensional vector, the three filters will not predict and remove a volume of random constant amplitude lines (spaghetti model). In that sense the three filters are a restriction of the general three-dimensional filter. Symbolically, the three fold prediction-error filtering amounts to

$$\epsilon(\mathbf{x}) = (A_x, A_y, A_z)^T g(\mathbf{x}), \tag{15}$$

where g is the original wavefield, A_x, A_y and A_z are the individual filter operators, and ϵ is the prediction-error vector.

The combination of the three filters generalizes the finite difference implementation of the cross-product operator of section *Plane wave misfit*. When trained at an identical plane-wave volume, the prediction-error filters and the finite difference approximations of the crossproduct operator are essentially identical. When encountering two competing plane-wave components separated by a fault, the cross-product operator estimates an average dip that leads to a large residual of the crossproduct everywhere. Similarly, the three prediction-error filters adjust their coefficients to reject both plane waves. Since the plane waves are not superposed, however, the prediction error is significant everywhere.

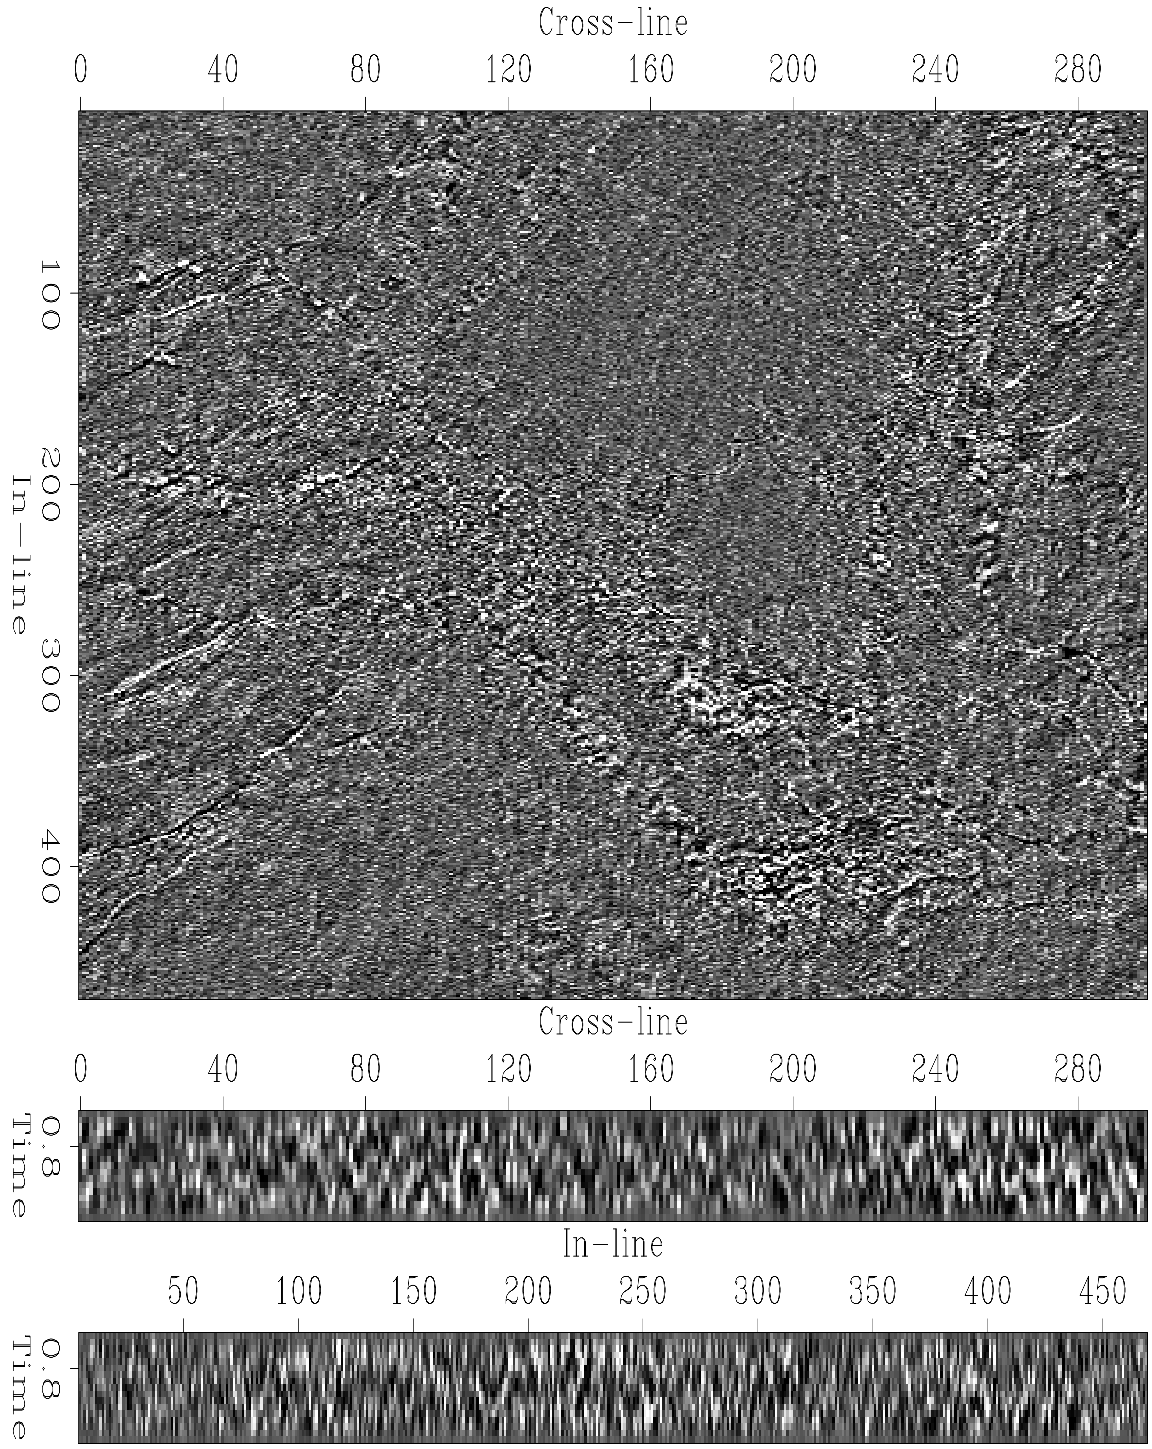

Figure 37: 3-D prediction-error of salt image. The attribute indicates the image's faults. In many locations, however, the attribute's noise obscures the fault features. gulfFoltTot1Pef3d [CR]

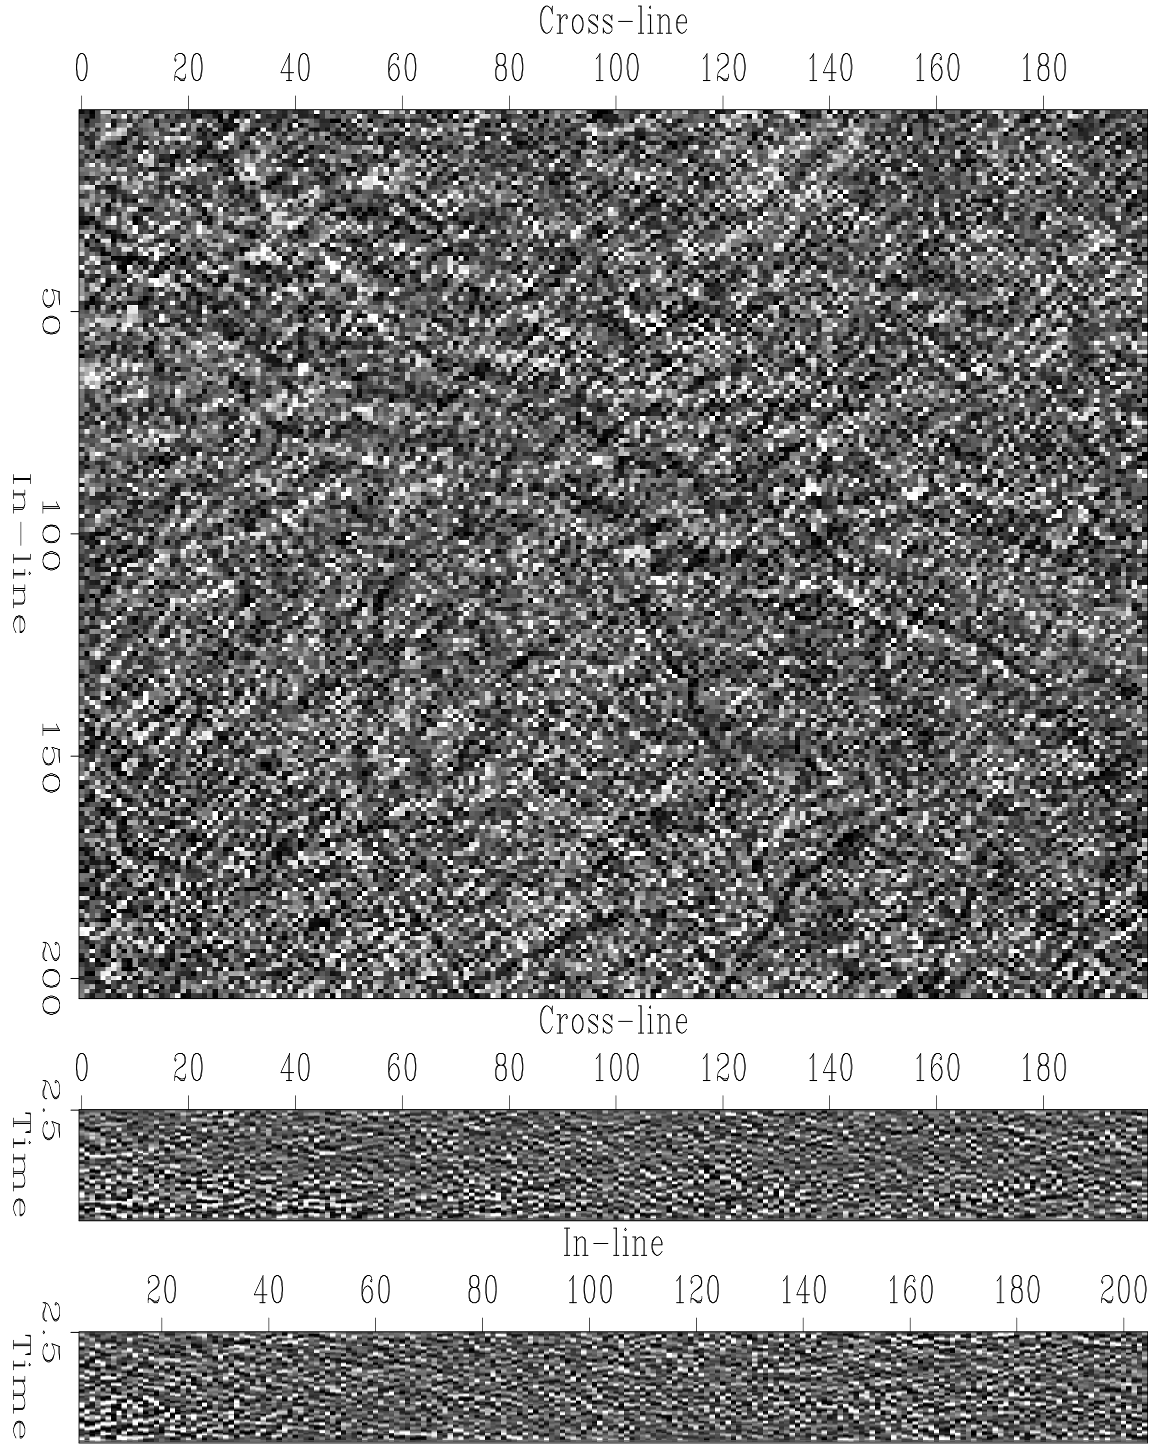

Figure 38: 3-D prediction-error of horst image. Among the noise, a pattern of linear events indicate the location of discontinuities to a careful observer. Unfortunately, the attribute image lacks clarity. nseaFoltTot1Pef3d [CR]

Results

Figure 40 shows the prediction error of the three two-dimensional filters applied to three copies of the synthetic test case. In all three images, the fault is delineated, but blurred. In patches away from the fault, the filters encounter and remove single plane wave events. In the patches that straddle the discontinuity, the filters encounter two plane wave volumes of distinct dip. Since the filters are limited to the removal of a single plane wave, each filter generates a large prediction error in those patches. The three-dimensional vector-field format is, of course, unsuitable for human interpretation. The next sections suggest various methods that unify the vector field into a single scalar subsurface map.

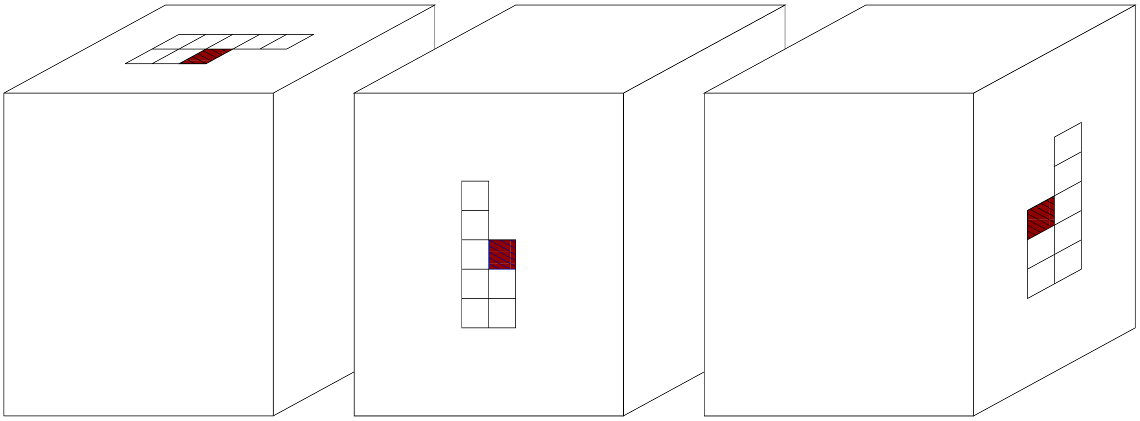

Figure 39: Three two-dimensional prediction-error filters. Each of three copies of the synthetic test case is filtered by one of the two-dimensional prediction-error filters of Figure 39. All three prediction-error volumes are zero, if and only if the input image is a single plane-wave. rayab3Doper [NR]

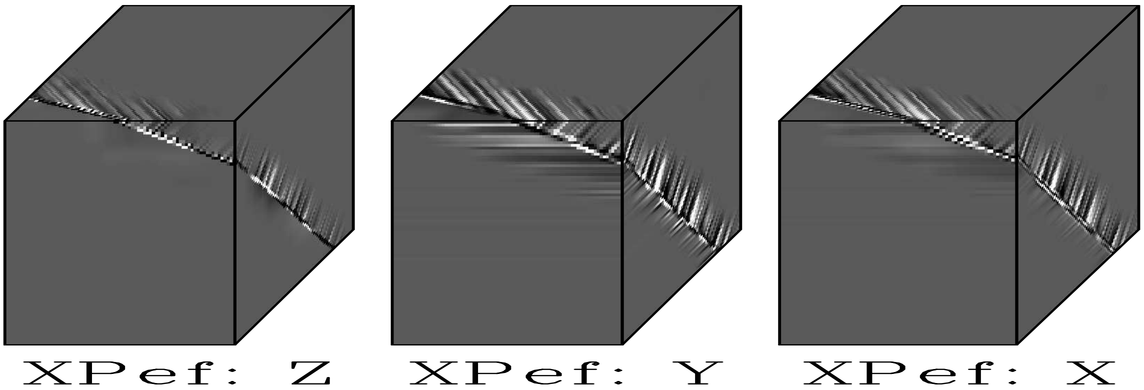

Figure 40: Three 2-D prediction-error of synthetic image. In all three volumes, the prediction error delineates the fault. The fault image is blurred, since the local filter is limited to removing only a single plane wave perfectly. The three-dimensional vector field is, of course, unsuitable for human interpretation and, therefore, requires further processing. zeroFoltXPef [ER]

NORM OF RESIDUAL

The norm of the prediction-error vector $\epsilon(\mathbf{x})$ of expression (15) collapses the three image to one:

$$|\epsilon(\mathbf{x})| = |(A_x, A_y, A_z)^T g(\mathbf{x})|.$$

Results

Figure 41 shows that the norm of the three prediction errors broadly delineates the fault in the synthetic case. Again the patches that contain the discontinuity show large prediction errors.

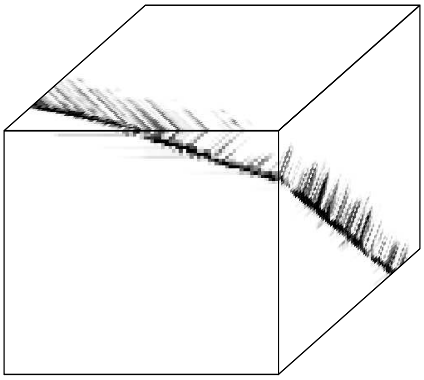

Figure 41: Magnitude of three 2-D prediction error of synthetic image. The fault is broadly outlined by patches of considerable prediction error. $\boxed{\text{zeroFoltXPefLN}}$ [ER]

The seismic test image of the Gulf salt dome indicates major fault lines as fuzzy zones of heightened prediction error (Figure 42). However, the overall image is noisy and the discontinuities are not pinpointed nor well contrasted from their surroundings. The sedimentary layering is completely suppressed. The salt gives rise to a region of low but noisy prediction error. Half of the pentagonoid region within the salt is clearly delineated; the other half is not. Additionally, the prediction-error attribute reveals a few features of the salt boundary that were obscured in other attribute images. For example, the long north-south fault at the eastern edge of the image. The North Sea horst image (not shown) similarly hints at its faults without differentiating and resolving them sufficiently. The three two-dimensional prediction error does not yield a convenient discontinuity image.

BACKPROJECTION OF RESIDUAL

An alternative method to combine the vector field $\epsilon(\mathbf{x})$ of expression (15) to a scalar field is the combination of the forward and adjoint of the prediction-error operator:

$$\epsilon(\mathbf{x}) = (A_x^T, A_y^T, A_z^T) \cdot (A_x, A_y, A_z)^T g(\mathbf{x}).$$

The space of the output ϵ is equal the space of the input g. Unfortunately, I do not know a insightful interpretation of the procedure. If prediction-error filter is similar to a derivative

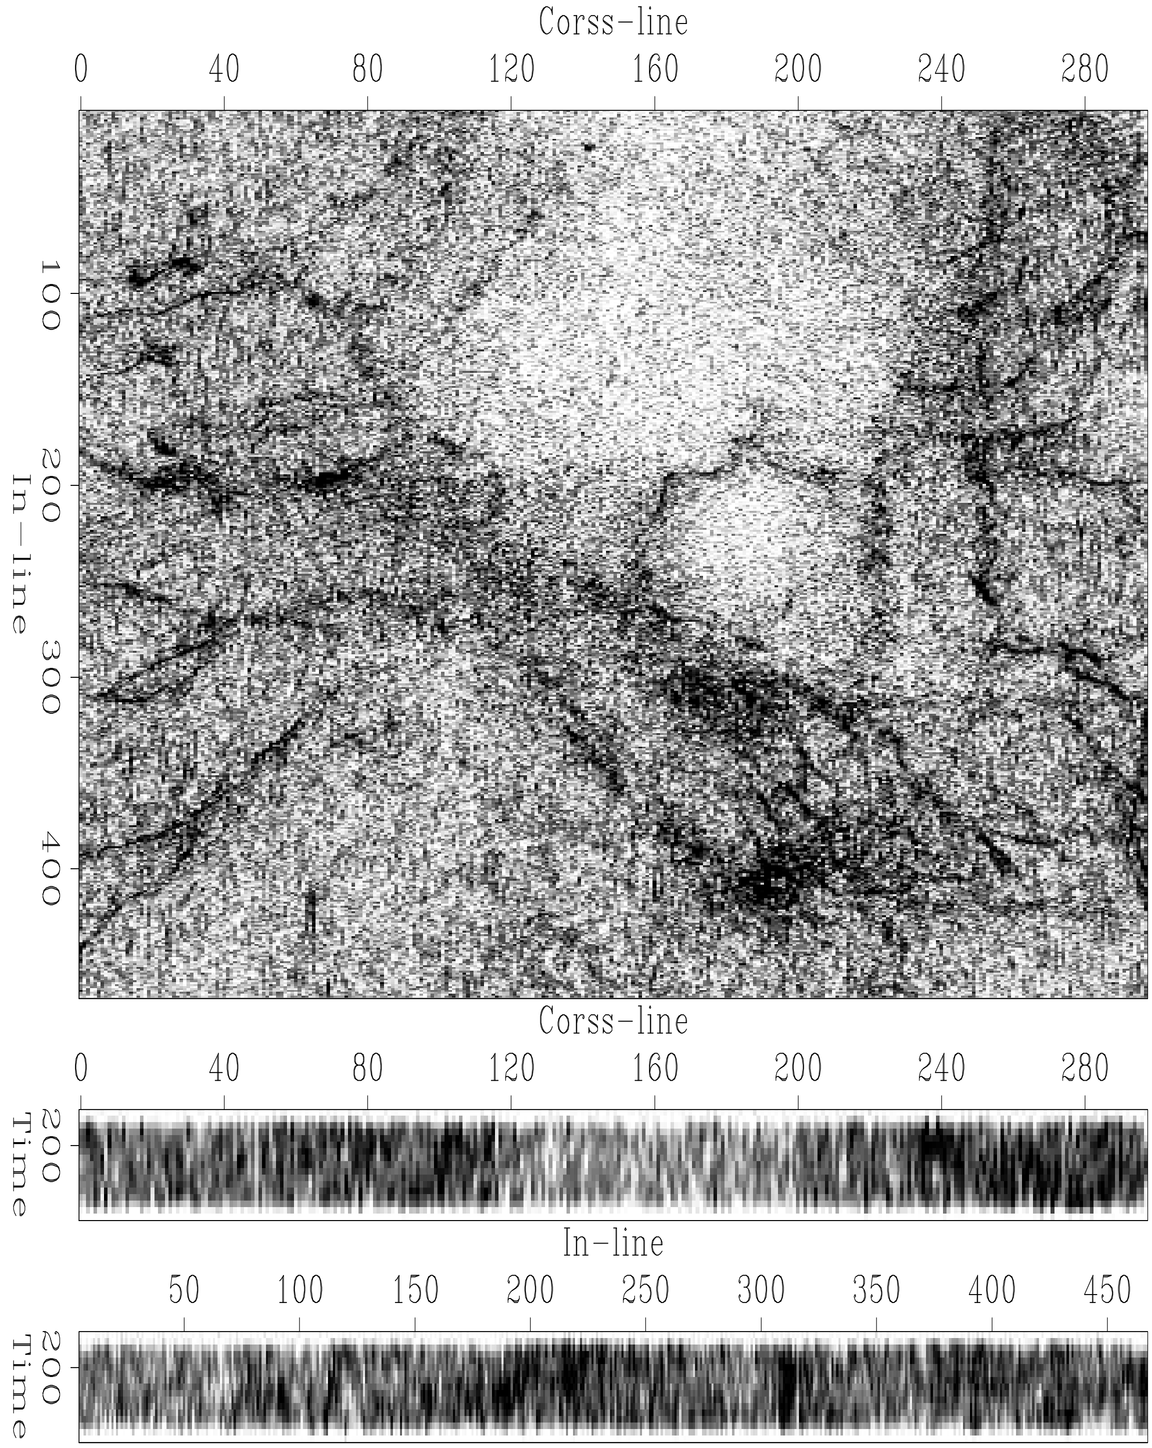

Figure 42: Magnitude of three 2-D prediction error of salt image. Broad, fuzzy zones indicate the image's major faults among a noisy background. gulfFoltTotXPefLN [CR]

operator, the combination of forward and adjoint corresponds to a sequence of two derivative operations.

Results

Figure 43 shows the discontinuity map of the synthetic image computed by the combination of the forward and adjoint of the prediction-error operator. The plane-wave patches are zeroed. The fault is delineated. The residual due to the nonstationarity of the discontinuous patches blurs the discontinuity. However, the operator seems to pinpoint the fault better than the other operators.

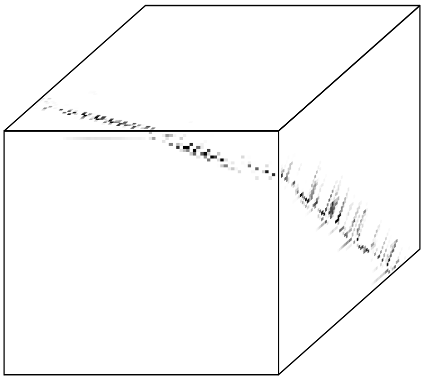

Figure 43: Forward and adjoint of three 2-D prediction-error filters applied to synthetic test case. The output delineates the fault, but surrounds it with some blur. zeroFoltXPefAA [ER]

The back-projected prediction error suppresses most of the wavefield character of the Gulf salt dome image (Figure 44. Some ringing within the salt body are due to the salt's weak original internal reflections. The salt's radial faults are mostly delineated. Especially, the fault marked R in the original image 8 appears in great details. Unfortunately, the image is noisy and other faults are slightly obscured. The salt body itself clearly delineates the pentagonoid region. The salt boundary is a broad zone of increased prediction error. The image clearly suffers from the visible presence of the north-south streaks that might be acquisition footprint. In summary, the discontinuity attribute shows many subtle details, however, it is too noisy to be quickly and reliably interpreted.

Figure 45 fails to indicate the location of the major faults, which exist in the original image. The overall wavefield character of the image is replaced by a white noise character. Despite temporal smoothing, the discontinuity image only hints at some of its major faults.

DISCUSSION

Contrary to my initial theoretical expectation, prediction error is not a reliable discontinuity attribute. In particular, I expected the combination of the three two-dimensional filters to remove any plane wave and to yield a sharp prediction error where they encountered anything but a plane wave. I was naively wrong on two accounts.

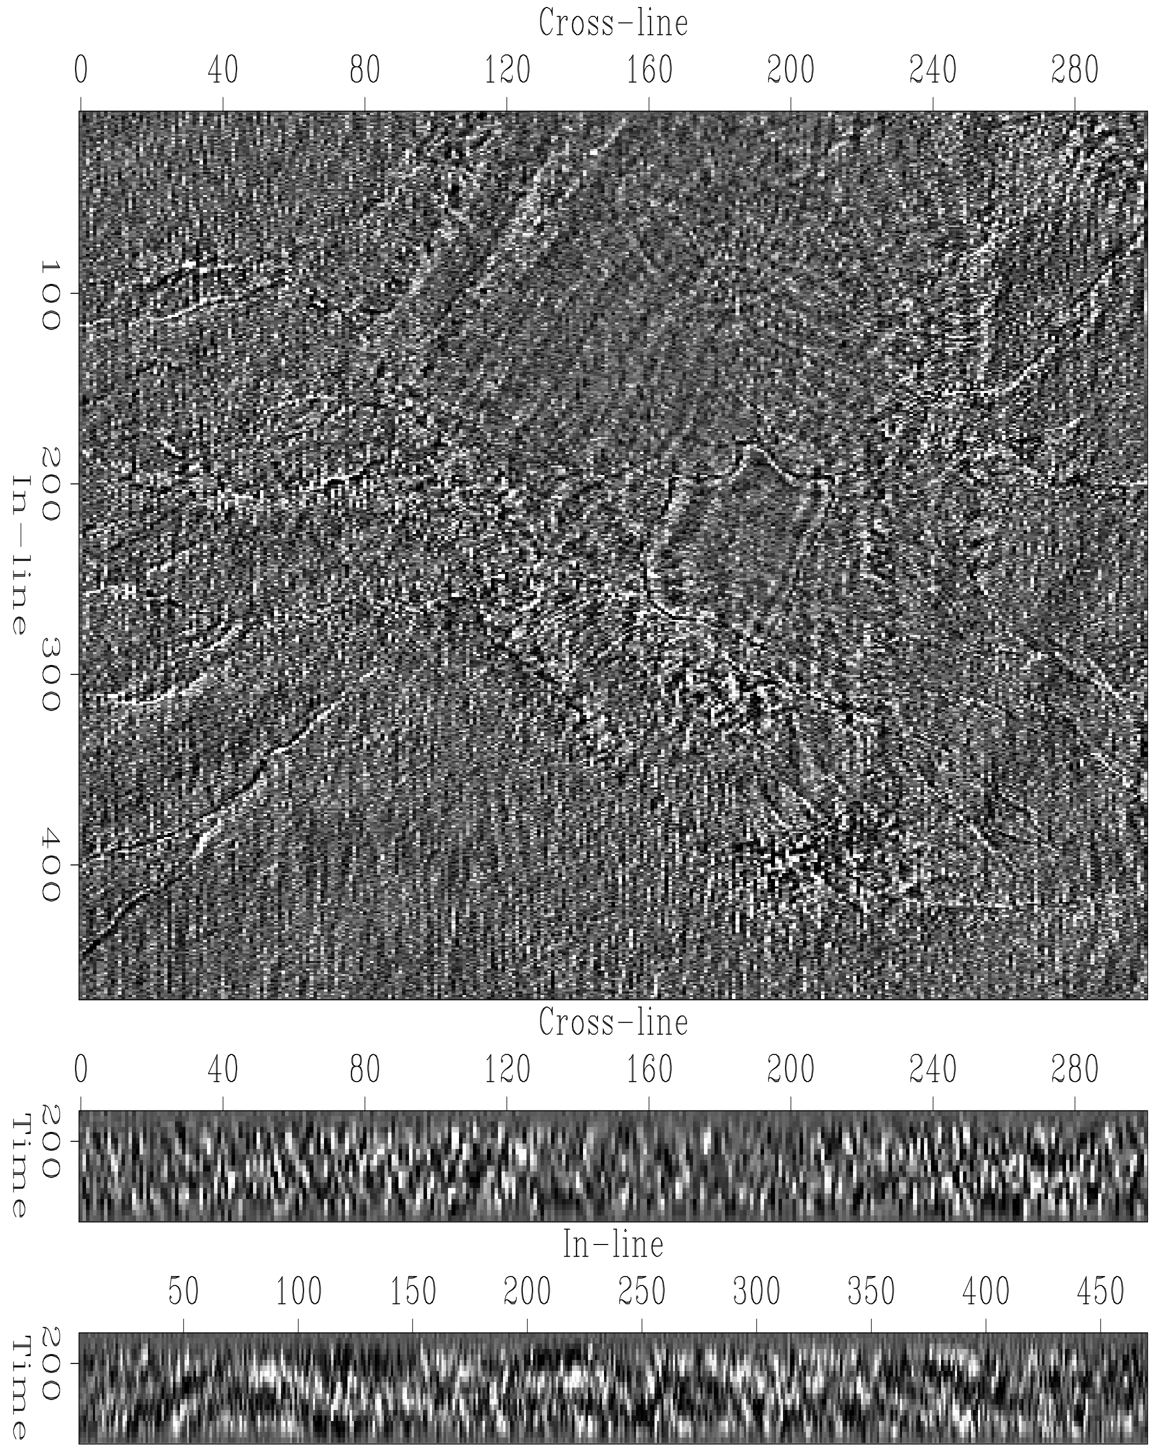

Figure 44: Forward and adjoint of three 2-D prediction error of salt image. The image delineates the radiating faults, however, the attributes noise level prevents the quick and reliable interpretation of the image. gulfFoltTotXPefAA [CR]

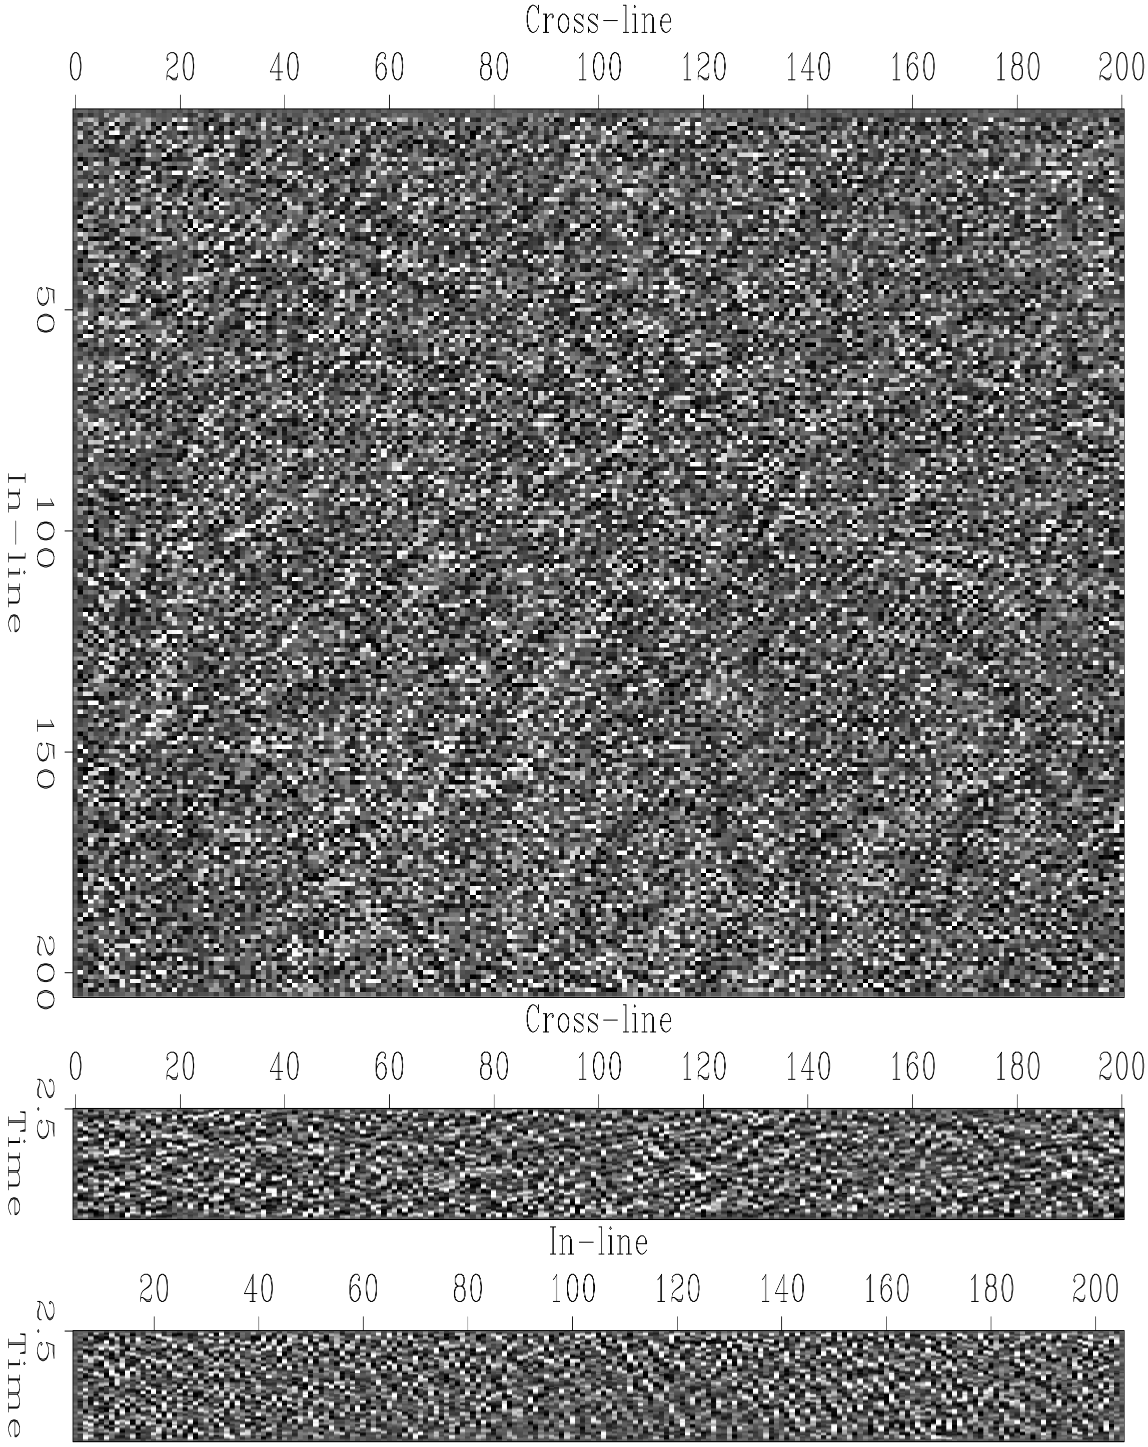

Figure 45: Forward and adjoint of three 2-D prediction-error filters. The discontinuity image appears noisy and fails to delineate any image features clearly. nseaFoltTotXPefAA [CR]

First, I failed to realize that prediction error indicates patches that contain a discontinuity but does not pinpoint the discontinuity within the patch. I disregarded the fundamental difference between the prediction of two superposed plane waves and two adjacent plane waves. Patching breaks the initial nonstationary image into small patches. Patches that contain discontinuities are, however, still nonstationary since they contain two adjacent plane-wave volumes. In general, a single prediction-error filter cannot perfectly predict and remove the statistically distinct image regions separated by the discontinuity. Instead, the filter generates considerable error amplitudes anywhere in the patch. The resolution of a prediction-error map is consequently its patch size. For example, the blurred fault of the synthetic image cases of Figures 28 and 41 illustrate the resolution loss. In this figure the patch size is about an eighth of the image size. Similarly, the plane-wave misfit of the previous chapter detects nonstationary patches that are not approximated by a single plane-wave.

Second, I incorrectly assumed that faults were sharp pixel-to-pixel boundaries, as I stylized them in the synthetic example. Instead, image processing is often unable to deliver sharp images of discontinuities. Such smooth transition zones do not cause large prediction errors. In particular, if the size of a fault zone is similar to the size of patch, the filter will be able to predict and remove components of the fault.

The patch size has to be large enough for the reliable estimation of of the prediction-error filter coefficients. On the other hand, the patch has to be small enough to contain stationary data in most cases and to yield a high-resolution output image. For the various examples, a typical patch size is $(nx, ny, nz) = (3, 3, 10)$.

In the various test cases, the norm of the three 2-D prediction-error filters broadly delineates seismic discontinuities as fuzzy lines (Figure 42). In contrast, the combination of the forward and adjoint of the prediction-error filter pinpoints the discontinuities better. It even successfully highlights features within the central salt region (Figure 44).

In general, I found prediction error an unreliable and noisy discontinuity attribute. Prediction-error fails to detect smooth fault zones. However, prediction error delineates sharp faults. The discontinuity maps are intricate and reveal details in regions with little amplitude contrast. However, prediction-error maps are obscured by noise and are neither quickly and nor easily interpreted.

ACKNOWLEDGMENTS

Jon Claerbout previously developed a two-dimensional and three-dimensional prediction-error scheme for the annihilation of plane waves. I used many of his excellent Fortran programs (or modified versions of them), until I changed to the Java programming language.

REFERENCES

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

1058.

Bleistein, N., 1984, Mathematical methods for wave phenomena: Academic Press.

Castleman, K. R., 1996, Digital image processing: Prentice-Hall, Inc., New Jersey.

Claerbout, J. F., 1992, Earth Soundings Analysis: Processing Versus Inversion: Blackwell Scientific Publications.

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.

Jain, A. K., 1989, Fundamentals of digital image processing: Prentice-Hall, Inc., New Jersey.

Neidell, N. S., and Taner, M. T., 1971, Semblance and other coherency measures for multichannel data: Geophysics, **36**, no. 3, 482–497.

Schwab, M., and Claerbout, J., 1995, The interpolation of a $3-d$ data set by a pair of $2-d$ filters: SEP–**84**, 271–278.

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

Schwab, M., 1997, Pre-whitening and coherency filtering: SEP–**94**, 189–204.

Sonneland, L., Barkved, O., and Hagenes, O., 1990, Construction of reservoir maps from seismic classifier maps: 60th Annual Internat. Mtg., Soc. Expl. Geophys., Expanded Abstracts, 241–244.

Sonneland, L., 1983, Computer aided interpretation of seismic data: 53rd Annual Internat. Mtg., Soc. Expl. Geophys., Expanded Abstracts, Session:S20.4.

Taner, M. T., Koehler, F., and Sheriff, R. E., 1979, See erratum. Volume 44, Issue 11, Page 1896.Complex seismic trace analysis: Geophysics, **44**, no. 6, 1041–1063.

Tarantola, A., 1987, Inverse problem theory: Elsevier.

Weglein, A. B., 1989, What can inverse scattering really do for you today?: Geophysical Inversion Workshop, Siam, Geophysical inversion: proceedings, 20–45.