

Detection and Extraction of Fault Surfaces in 3-D Seismic Data

Israel Cohen^{1*}, Nicholas Coult² and Anthony A. Vassiliou²

¹Department of Electrical Engineering, Technion - Israel Institute of Technology, Haifa, Israel
²GeoEnergy, Inc., 3000 Wilcrest Drive, Houston, Texas

Abstract

In this paper, we propose an efficient method for detecting and extracting fault surfaces in 3-D seismic volumes. The seismic data is transformed into a volume of Local Fault Extraction (LFE) estimates, representing the likelihood that a given point lies on a fault surface. We partition the fault surfaces into relatively small linear portions, which are identified by analyzing tilted and rotated subvolumes throughout the region of interest. Directional filtering and thresholding further enhance the seismic discontinuities attributable to fault surfaces. Subsequently, the volume of LFE estimates is skeletonized and individual faults surfaces are extracted and labeled in a decreasing order of their size. The ultimate result obtained by the proposed procedure provides a visual and semantic representation of a set of well-defined, cleanly separated, one-pixel-thick labeled fault surfaces, which is readily usable for seismic interpretation.

*Corresponding author. E-mail: icohen@ee.technion.ac.il; Tel.: +972 4 8294731; Fax: +972 4 8295757

INTRODUCTION

Fault surfaces are common subterranean structures, associated with displacements or offsets of subsurface layers. Their consistent and reliable detection in 3-D seismic data provides an interpreter with very powerful means to quickly visualize and map complex geological structures.

A common tool facilitating structural and stratigraphic interpretation is the coherency cube, originated by Bahorich and Farmer (1995, 1996). It is calculated from the seismic data using a coherency measure that quantifies the seismic discontinuity at each point. Discontinuities attributable to fault surfaces include dip, azimuth, and offset changes of seismic reflectors, and waveform and amplitude variations caused by defocusing. Such discontinuities appear on coherence slices as incoherent linear or curved features (*e.g.*, Marfurt *et al.* 1999; Gersztenkorn *et al.* 1999; Neff *et al.* 2000; Lees 1999).

The most acceptable coherence measures are based on cross correlation (Bahorich and Farmer 1995), semblance (Marfurt *et al.* 1998), or eigenstructure (Gersztenkorn and Marfurt 1996a, 1996b; Kirilin 1992) techniques. These methods typically suffer from a lack of robustness, especially when dealing with noisy data (Marfurt *et al.* 1999; Gersztenkorn and Marfurt 1999). Recently, we have introduced a multiscale analysis method for the estimation of seismic coherency, which is both robust to noise and computationally efficient (Cohen and Coifman 2002). It involves another measure, namely the *Local Structural Entropy* (LSE), which evaluates the dissimilarity of subvolumes enclosing a given analysis point. Dealing with subvolumes, rather than individual traces, leads to robustness, while avoiding the expensive computations of eigenstructure-based large covariance matrices and eigenvalues.

A major drawback of coherency-based fault analysis is that seismic discontinuities may also result from geological features, which are unrelated to faults. Furthermore, creating a consistent geological interpretation from large 3-D seismic data volumes often requires manual intervention,

which is time-consuming, tedious and imprecise.

In this paper, we propose a robust and computationally efficient method for the extraction of fault surfaces in 3-D seismic volumes. The seismic data is transformed into a volume of Local Fault Extraction (LFE) estimates, which provides the interpreter with a much clearer visual indication of the fault surfaces. The LFE estimate at a given analysis point is obtained by the following procedure. First, a 3-D analysis cube, tilted and rotated about the analysis point, is selected by the interpreter. The analysis cube moves throughout the seismic volume and outputs for each point a measure of Normalized Differential Entropy (NDE). The NDE value represents the likelihood of a fault surface, having similar dip and azimuth as of the analysis cube, to intersect with the analysis point. Subsequently, the local average of the NDE is removed, and portions of fault surfaces, approximately aligned with the analysis cube, are extracted by directional filtering. The filtered NDE coefficients are thresholded, and filtered back to produce directional LFE volumes. Next, the LFE attribute is given by the maximal directional LFE, over the presumably tested set of dips and azimuths. This approximately gathers the significant portions of the fault surfaces into smooth larger surfaces. Finally, the results are skeletonized, and individual faults surfaces are extracted and labeled in a decreasing order of their size. A comparison of the LFE volume with the Local Structural Entropy attribute (Cohen and Coifman 2002) shows that the LFE method provides a more reliable and convenient tool for detecting and extracting fault surfaces.

In the following section, we present the fault extraction algorithm. We begin with a local estimate of faults. Next, we take these faults and enhance the contrast, which facilitates the analysis of regions that contain dipping layers or are highly discontinuous. Then we extract the portions of fault surfaces by directional filtering, and construct fault surfaces out of these portions. Finally, we demonstrate the applicability of the proposed algorithm on real data.

THE LOCAL FAULT EXTRACTION

In this section we describe the basic components forming the proposed fault extraction algorithm.

Normalized differential entropy

We begin by subdividing the original three-dimensional seismic volume into individual 3-D data analysis volumes of $N \times M \times P$ dimension. A typical analysis volume is 41 values by 11 values by 6 values ($41 \times 11 \times 6$). The analysis volume is defined by the length L_1 along the major axis, the length $2L_2 + 1$ along the minor axis, time duration of N samples, azimuth ϕ (rotation angle around the in-line axis), and tilt θ from the vertical axis (see Figure 1).

The analysis volume is broken into two subvolumes (for instance, for a $41 \times 11 \times 6$ analysis cube, each of the two subvolumes is $41 \times 11 \times 3$ for a total of 1353 values) which are rotated and tilted around a central analysis point $p = (t, x, y)$. Subsequently, the samples within the respective subvolumes are rearranged in a consistent manner into two column vectors $\mathbf{v}_{1,p}(\theta, \phi)$ and $\mathbf{v}_{2,p}(\theta, \phi)$. It is worthwhile noting that the subvolumes in Figure 1 have horizontal top and bottom surfaces, which is particularly preferable when the subsurface layers are horizontal or close to horizontal. However, in case of dipping layers, the top and bottom surfaces can be elevated in parallel to the dominant local seismic orientation within the analysis cube (Cohen, Vassiliou and Coult, 2004).

By rearranging the seismic data within the analysis cube in such a way, the resulting column vectors can be directly used for edge detection. That is, the fault surface edges defined by a combination of different position of the central analysis point p , the analysis cube size, the rotation ϕ and the tilt θ are computed from the pairs of the corresponding two column vectors. The computations, which are carried out through a normalized version of the Prewitt edge detection filter (Jain, 1989; Lipkin and Rosenfeld, 1970; Luo *et al.*, 1996), are designed to capture the edges

and breaks visible or invisible, generated by subterranean faults in the seismic data.

A measure $\mathcal{N}_p(\theta, \phi)$, referred to as normalized differential entropy (NDE), is computed at each analysis point p as a normalized version of the Prewitt filter

$$\mathcal{N}_p(\theta, \phi) = \frac{\|\mathbf{v}_{1,p}(\theta, \phi) - \mathbf{v}_{2,p}(\theta, \phi)\|_{\mathbf{p}}}{\|\mathbf{v}_{1,p}(\theta, \phi)\|_{\mathbf{p}} + \|\mathbf{v}_{2,p}(\theta, \phi)\|_{\mathbf{p}}}, \quad (1)$$

where $\|\cdot\|_{\mathbf{p}}$ is the $\ell_{\mathbf{p}}$ norm (typically we use the ℓ_1 norm, and the metric defined in (1) can be referred to as a normalized Manhattan distance). We may observe the following: First, the range of $\mathcal{N}_p(\theta, \phi)$ is $[0, 1]$, where the minimal value (zero) is obtained when the two subvolumes are perfectly correlated without an offset of seismic layers (*i.e.*, $\mathbf{v}_{1,p}(\theta, \phi) = \mathbf{v}_{2,p}(\theta, \phi)$), and the maximal value of $\mathcal{N}_p(\theta, \phi)$ is obtained for maximally offset correlated subvolumes (*i.e.*, $\mathbf{v}_{1,p}(\theta, \phi) = -\mathbf{v}_{2,p}(\theta, \phi)$). Second, if the subsurface layers have orientation mostly horizontal with no significant lateral elastic impedance contrasts, then we have an NDE measure which produces an excellent edge detector and indicator of the fault surfaces and at the same time a very poor indicator of the subsurface layer interfaces. Third, even when the dominant layer orientation is not horizontal, by using top and bottom analysis cube surfaces parallel to the dominant local 3-D seismic orientation (Cohen, Vassiliou and Coult, 2004), we still obtain a very good 3-D edge detector and indicator of the fault surfaces and at the same time a very poor indicator of the subsurface layer interfaces. Fourth, the data structure employed for the NDE and its 3-D edge detection computational architecture is completely different than the data structures and the computational setup in the 3-D seismic coherence and variance methods (*e.g.*, Marfurt *et al.*, 1999, Bahorich and Farmer, 1996, Gersztenkorn *et al.*, 1999). Fifth, when the top and bottom analysis cube surfaces are horizontal, there is a very significant overlapping of the analysis cubes for any dip and azimuth combination, and accordingly the computation of $\mathcal{N}_p(\theta, \phi)$ and $\mathcal{N}_{p'}(\theta', \phi')$ for nearby analysis points p and p' involves many repeatable difference computations. Consequently, all of the possible differences for the NDE computations are computed upfront for all the combinations of the dip and azimuth used, which yields significant acceleration

in the NDE computations. We have found from practical experience that a set of 12 dips from 35 to 35 degrees with a 5 degree increment excluding the dips of 5, 0, 5 degrees, and a set of 12-15 uniformly-spaced azimuth values are sufficient for a $41 \times 6 \times 7$ NDE analysis cube. By using the upfront NDE local difference computations for all the combinations of dip and azimuth we obtain a computational speedup of about a factor of 20, which is extremely important for efficient computation and extraction of the fault surfaces.

Contrast enhancement

The second step of the algorithm is contrast enhancement, applied to single NDE volumes, *i.e.* per postulated orientation. Fault surfaces having dips and azimuths about the same dip and azimuth of the analysis cube are distinguished by higher NDE values, compared to the local average NDE values. Accordingly, we apply a contrast enhancement filtering to single NDE volumes, *i.e.* per postulated orientation, and set to zero negative values. This contrast enhancement facilitates the analysis of regions that contain dipping layers or are highly discontinuous.

The contrast enhancement can be efficiently implemented by using a discrete “Mexican hat” function, which corresponds to local average subtraction:

$$f(n) = C(1 - n^2T^2) \exp(-n^2T^2/2), \quad (2)$$

where T is the sampling period of the Mexican hat function, n is an integer, and C is a normalization constant such that $\sum_{n=-\infty}^{\infty} |f(n)| = 2$. We use a finite length filter ($-4.5 \leq nT \leq 4.5$), containing odd number of uniformly spaced coefficients. In general, the selected filter length depends on the size of the analysis cube and the “thickness” of the fault surfaces. The filtered NDE by the Mexican hat is computed as

$$\bar{\mathcal{N}}_p(\theta, \phi) = g_p(\theta, \phi) * \mathcal{N}_p(\theta, \phi) = \sum_{p'} g_{p-p'}(\theta, \phi) \mathcal{N}_{p'}(\theta, \phi), \quad (3)$$

where $g_p(\theta, \phi)$ is a rotated 3-D version of f around a central analysis point $p = (t, x, y)$, such that its main axis is perpendicular to the slabs of the analysis cube (*i.e.*, rotated by azimuth ϕ around the in-line axis, and tilted by θ from the vertical axis). The contrast enhanced NDE is given by

$$\hat{\mathcal{N}}_p(\theta, \phi) = \max \{ \bar{\mathcal{N}}_p(\theta, \phi), 0 \} . \quad (4)$$

Directional LFE

The third step of the Fault Mapping System utilizes 3-D directional filtering on single contrast enhanced NDE volumes, *i.e.* per postulated orientation. The directional filtering enhances the portions of fault surfaces that are approximately aligned with the analysis cube.

The directional filter, denoted by $h_p(\theta + \alpha, \phi)$, is a 3-D ellipsoid, tilted by $\theta + \alpha$ with respect to the time axis, rotated by ϕ with respect to the in-line axis, and normalized by $\sum_p h_p(\theta, \phi) = 1$. Its dimensions, selected by the interpreter, control the minimal dimensions of the detected subsurfaces. The maximum value of α is determined by the tilt increment Δ_θ ($|\alpha| < \Delta_\theta/2$). The implementation uses a 3-D pancake-like shaped Hann window, aligned parallel to the division in the analysis cube. A possible set of dimension values for this 3-D pancake-like window are 61 samples at its major axis and 3 samples at its minor axes. A possible value for the tilt increment is $\Delta_\theta = 5^\circ$, and the relative tilt of the directional filter α can be restricted to values in the set $\{-2^\circ, 0, 2^\circ\}$. Alternatively, a smaller tilt increment could be used and the relative tilt could be discarded (*i.e.*, $\alpha \equiv 0$); however, the above formulation has been found to be computationally more efficient.

Directional filtering of the contrast enhanced NDE yields

$$C_p(\theta + \alpha, \phi) = \sum_{p'} h_{p-p'}(\theta + \alpha, \phi) \hat{\mathcal{N}}_{p'}(\theta, \phi), \quad (5)$$

where the summation in (5) is over the set of points $p' = (t', x', y')$ which are in the proximity of $p = (t, x, y)$ (*i.e.*, $p - p'$ is in the support of $h_{p-p'}(\theta + \alpha, \phi)$). The resulting coefficients are

thresholded by δ ($0 < \delta < 1$),

$$\tilde{C}_{p'}(\theta + \alpha, \phi) = \begin{cases} C_{p'}(\theta + \alpha, \phi), & \text{if } C_{p'}(\theta + \alpha, \phi) \geq \delta \\ 0, & \text{otherwise,} \end{cases} \quad (6)$$

and then filtered back to produce the directional LFE values, determined by

$$\mathcal{L}_p(\theta, \phi) = \sum_{p', \alpha} \tilde{C}_{p'}(\theta + \alpha, \phi) h_{p-p'}(\theta + \alpha, \phi). \quad (7)$$

The directional LFE volumes contain significant portions of fault surfaces, characterized by roughly the same dip and azimuth orientations as those of the analysis cube. Subsequently, we keep at each point of the 3-D volume the maximum directional LFE value over the tested set of tilts and azimuths. Specifically, the LFE attribute at the analysis point p is obtained by

$$\hat{\mathcal{L}}_p = \max_{\theta, \phi} \{\mathcal{L}_p(\theta, \phi)\}. \quad (8)$$

Equipped with the tilt and azimuth arguments, $\theta(p)$ and $\phi(p)$, which yield maximum in (8) for each particular analysis point p , the LFE volume is further enhanced by 3-D skeletonization and 3-D surface separation. The directional LFE performs best on 3-D seismic volumes that do not have fault shadows. Extensive presence of fault shadows can generate small fault surfaces which are crossing true faults in an X-pattern.

Skeletonization

The fourth step of the Fault Mapping System involves a skeletonization step, which is very important in both “filtering out” very small faults or fault-like features and for separating the different fault surfaces. Skeletonization algorithms have been in use in image processing for a significant time (*e.g.*, Blum 1967, Stentiford and Mortimer 1983, Zhang and Fu 1984). Here, we present a skeletonization algorithm for 3-D data which is particularly designed for the results of the filtering step outlined above.

The skeletonization algorithm performs the following computations on the results of the filtering step:

1. For each horizontal slice ‘t’, a 2-D binarization and skeletonization is performed using an “adaptive image binarization and skeletonization” algorithm (see Appendix A).
2. For each vertical slice ‘x’ perform 2-D skeletonization, and stretch the skeletons, as described in the “adaptive image binarization and skeletonization” algorithm (see Appendix A).
3. Repeat Step 2 for vertical slice ‘y’, horizontal slice ‘t’, vertical slice ‘x’, and so on, until there is no change in the result, or reached to a prespecified limit of iterations. Each iteration further stretches the skeletons, so that gaps in the fault surfaces are bridged and the resulting surfaces become more complete.

Labeling of individual fault surfaces

The final step of the Fault Mapping System involves separating the individual fault surfaces, and labeling the resulting objects in a decreasing order of their size.

For each separate azimuth, we extract the data associated with that azimuth and concatenate the extracted data along the fourth dimension, ordered according to the azimuth value, thus producing a 4-D data that contains the azimuth information as well. Subsequently, we label the newly produced 4-D data and re-create a 3-D labeled data by taking the maximum along the fourth dimension. This will create a labeled version of the 3-D data, such that objects are distinguished by their azimuths as well, and therefore differently labeled if their azimuths are not close one to another. Next, we identify objects whose size is above a certain threshold, and re-label these objects in a decreasing order of their size (the remaining objects are removed).

RESULTS

In this section we use a real data example to demonstrate the applicability of the LFE algorithm, and to illustrate its execution. The data is decimated in both time and space. The time interval is 4 ms, in-line and cross-line trace spacing is 25 m. A small subvolume of 401 x 401 traces and 1001 samples is used for demonstration (10 km in-line and cross-line distances and 4 s duration).

Figure 2 shows vertical cross-sections and horizontal slices through the seismic data. The corresponding vertical cross-sections and horizontal slices through the skeletonized LFE volume are displayed in Figure 3. For the present example, we used a 3-D pancake-like shaped Hann filter, whose dimensions are 61 samples at its major axis and 3 samples at its minor axes. The tilt increment is $\Delta_\theta = 5^\circ$, the azimuth increment is $\Delta_\phi = 45^\circ$, and the relative tilt of the directional filter α is restricted to three values $\{-2^\circ, 0, 2^\circ\}$. The filtered NDE coefficients are thresholded by $\delta = 0.12$, and filtered back to produce the directional LFE volumes. The skeletonized LFE volume (Figure 3) contains fault surfaces which are consistent with the presumed model (*i.e.*, the dimensions of the analysis cube, set of dips and azimuths, directional filter, threshold, *etc.*). Due to the existence of fault shadows in the prestack time migrated volume, small scale X-pattern faults are generated. Hence it is generally preferable to apply LFE on prestack migrated data with minimal amount of fault shadows. Figure 4 shows the final result of the Fault Mapping System, which includes labeling of distinct skeletons in a decreasing order of their size.

For comparison with the Local Structural Entropy attribute (Cohen and Coifman 2002), we plot in Figure 5 vertical cross-sections and horizontal slices through the Local Structural Entropy volume, obtained by using analysis cubes of 6×6 traces and 21 samples. Values are mapped to shades of gray, where darker shades indicate greater discontinuity. A comparison of Figures 5 and 3 shows that the proposed method provides a more reliable and convenient tool for the extraction

and separation of fault surfaces. Furthermore, faults along strike which are difficult to see in the Local Structural Entropy volume, are very clear in the skeletonized LFE volume.

CONCLUSION

A major component of structural interpretation is the identification, location, and extraction of individual fault surfaces. Fault surfaces are very important in hydrocarbon exploration since they are directly related to hydrocarbon accumulation and to hydrocarbon flow paths. Extraction of individual fault surfaces from seismic data is a largely qualitative procedure and has thus been characterized by the need for careful human data interpretation.

In recent years, there has been progress in visualizing stratigraphic and structural discontinuities with the so-called coherence or variance methods, which look at the similarity or dissimilarity of a small number of neighboring traces to determine these discontinuities. The existing methods, however, have a limited precision, and are generally inefficient in isolating fine details, such as fault surfaces or sedimentary layer interfaces that may be only one pixel wide. The proposed algorithm facilitates the extraction, separation and labeling of the fault surfaces, by identifying portions of the surfaces and combining them into large distinct fault surfaces. Specifically, the portions of the fault surfaces are identified by analyzing tilted and rotated subvolumes throughout the region of interest. Applying 3-D contrast enhancement and directional filtering enables the analysis of regions that contain dipping layers or are highly discontinuous. Skeletonizing the results of the filtering operation, separating the individual fault surfaces, and finally labeling the skeletons in a decreasing order of their size yields a useful representation for a three-dimensional visualization of subterranean features.

Appendix A - Adaptive Image Binarization and Skeletonization

The steps of the adaptive image binarization and skeletonization algorithm are as follows:

1. Define two thresholds - low and high. The high threshold should be defined such that low intensity objects will be removed by the high threshold binarization and yet objects that are supposed to be kept in the image will not be completely erased by the high threshold binarization. The low threshold should be defined such that connectivity between two close objects will be gained and yet no unnecessary pixels will be marked as '1'.
2. Binarize the original image using the high threshold.
3. Skeletonize the binary image as follows:
 - (a) Using an 8-neighborhood notation as shown in Figure 6(a), let $N(p_1)$ denote the number of nonzero neighbors of the center point p_1 , and let $T(p_1)$ denote the number of 0 – 1 transitions in the ordered sequence $p_2, p_3, \dots, p_8, p_9, p_2$. Compute $N(p)$ and $T(p)$ for all points p in the image
 - (b) Flag type I border points in the image, where a point p_1 is flagged if the following conditions are satisfied:
 - i. $2 \leq N(p_1) \leq 6$;
 - ii. $T(p_1) = 1$;
 - iii. $p_5 = 0$ or $p_7 = 0$ or $(p_3 = 0$ and $p_9 = 0)$.
 - (c) Delete all type I border points.
 - (d) Flag type II border points in the image, where a point p_1 is flagged if the following conditions are satisfied:
 - i. $2 \leq N(p_1) \leq 6$;

- ii. $T(p_1) = 1$;
 - iii. $p_3 = 0$ or $p_9 = 0$ or $(p_5 = 0$ and $p_7 = 0)$.
- (e) Delete all type II border points.
- (f) Iterate deletion of type I and type II points through steps 3(b)-(e) until there are no more such border points in the image.
4. For each pixel in the skeletonized image that is marked as ‘1’ and has no neighbors marked as ‘1’ or has only one neighbor marked as ‘1’ (edge points) do the following:
- (a) For a pixel that has no neighbors marked as ‘1’, find the pixel with the maximum value among the 8-connected neighborhood of the corresponding pixel in the original image. For a pixel that has only one neighbor marked as ‘1’, find the pixel with the maximum value among the pixels in the original image that belong to the 8-connected neighborhood of the pixel and are “not too close” to the neighbor pixel marked as ‘1’. The locations of the pixels that are “not too close” to the neighbor pixel marked as ‘1’ are shown in Figure 6(b) (the left figure refers to the possibility that the neighbor marked as ‘1’ belongs to the 4-connected neighborhood and the right figure refers to the other possibility).
 - (b) If this maximum value is above the low threshold, mark the corresponding pixel in the skeletonized image as ‘1’. Otherwise try to find such a pixel (*i.e.*, one that is above the low threshold) in the 5×5 neighborhood in a similar manner. If such a pixel exists, mark the corresponding pixel in the skeletonized image as ‘1’ and also mark the appropriate pixel in the 3×3 neighborhood so connectivity will be retained.
 - (c) If the newly marked pixel (the outer pixel in the case of 5×5 neighborhood) has only one neighbor marked as ‘1’, go to step 4(a) in the algorithm for the newly marked pixel.

It is important to note that the skeletonized image is updated as necessary at the various steps of the algorithm, and every step in the algorithm uses the previously updated skeletonized image.

Furthermore, steps 4(a) through 4(c) in the algorithm can be implemented as a recursive function, which is recursively called as necessary in step 4(c) of the algorithm.

ACKNOWLEDGEMENTS

The authors thank Prof. Ronald R. Coifman of Yale University for valuable discussions and helpful suggestions. We also thank the Associate Editor, Prof. Kurt J. Marfurt, and the three anonymous reviewers who provided valuable comments which considerably improved the quality of this paper.

REFERENCES

- Bahorich, M. S., and S. L. Farmer, 1995, 3-D seismic discontinuity for faults and stratigraphic features: The Leading Edge, **14**, 1053–1058.
- Bahorich, M. S., and S. L. Farmer, 1996, Methods of seismic signal processing and exploration: US Patent 5 563 949.
- Blum, H., 1967, A transformation for extracting new descriptors of shape, in Weiant Whaten-Dunn, ed., Symposium models for speech and visual form: Cambridge, MA, MIT Press.
- Cohen, I., and R. R. Coifman, 2002, Local discontinuity measures for 3-D seismic data: Geophysics, **67**, 1933–1945.
- Cohen, I., A. Vassiliou and N. Coult, 2004, "Method and Apparatus for Seismic Feature Extraction", Patent number WO2004044615.
- Gersztenkorn, A., and K. J. Marfurt, 1996a, Coherence computations with eigenstructure: 58th Conf. and Tech. Exhibition, Eur. Assn. Geosci. Eng., Extended Abstracts, X031.
- Gersztenkorn, A., and K. J. Marfurt, 1996b, Eigenstructure based coherence computations: 66th Ann. Internat. Mtg., Soc. Expl. Geophys., Extended Abstracts, 328–331.

- Gersztenkorn, A., and K. J. Marfurt, 1999, Eigenstructure-based coherence computations as an aid to 3-D structural and stratigraphic mapping: *Geophysics*, **64**, 1468–1479.
- Gersztenkorn, A., J. Sharp, and K. J. Marfurt, 1999, Delineation of tectonic features offshore Trinidad using 3-D seismic coherence: *The Leading Edge*, **18**, 1000–1008.
- Jain, A. K., 1989, *Fundamentals of digital image processing*: Prentice Hall Inc.
- Kirlin, R. L., 1992, The relationship between semblance and eigenstructure velocity estimators: *Geophysics*, **57**, 1027–1033.
- Lees, J. A., 1999, Constructing faults from seed picks by voxel tracking: *The Leading Edge*, 338–340.
- Lipkin, B. S., and A. Rosenfeld (Eds.), 1970, *Picture processing and psychopictorics*: Academic Press Inc.
- Luo, Y., W. G. Higgs and W. S. Kowalik, 1996, Edge detection and stratigraphic analysis using 3D seismic data: 66th Annual Meeting, Expanded Abstracts, Society of Exploration Geophysicists, 324–327.
- Marfurt, K. J., R. L. Kirlin, S. L. Farmer, and M. S. Bahorich, 1998, 3-D seismic attributes using a semblance-based coherency algorithm: *Geophysics*, **63**, 1150–1165.
- Marfurt, K. J., V. Sudhaker, A. Gersztenkorn, K. D. Crawford, and S. E. Nissen, 1999, Coherency calculations in the presence of structural dip: *Geophysics*, **64**, 104–111.
- Neff, D. B., J. R. Grismore, and W. A. Lucas, 2000, Automated seismic fault detection and picking: US Patent 6 018 498.
- Stentiford, F. W. M. and R. G. Mortimer, 1983, Some new heuristics for thinning binary hand-printed characters for OCR: *IEEE Transactions on Systems, Man and Cybernetics*, **13**, 81–84.

Zhang, S. and K. S. Fu, 1984, A thinning algorithm for discrete binary images: International Conference on Computers and Applications, Beijing, China, 879–886.

Figure Captions

Figure 1: (a) A vertical cross-section and (b) horizontal slice illustrating the geometrical distribution of traces and samples used in the analysis cube. The analysis cube, consisting of two subvolumes, is centered about an analysis point $p = (t, x, y)$, and defined by length of major axis L_1 , length of minor axis $2L_2 + 1$, time duration N samples, azimuth ϕ around the in-line axis, and tilt θ from the vertical axis.

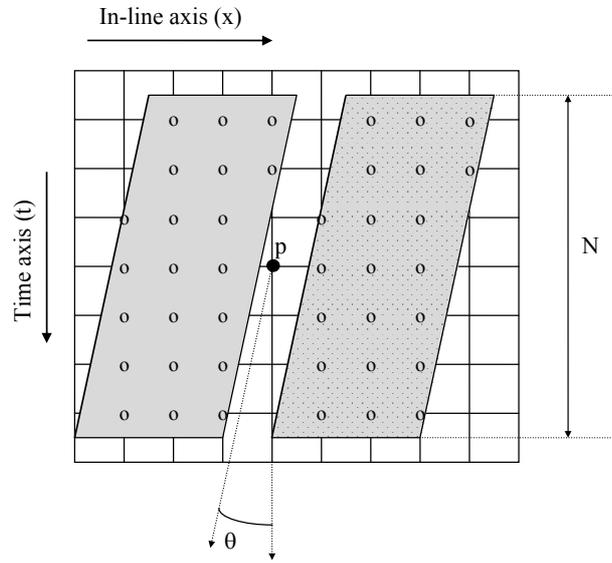
Figure 2: Vertical cross-sections and horizontal slices through the seismic data: (a) Vertical cross-section at $y = 1.3$ km; (b) vertical cross-section at $y = 8$ km; (c) horizontal slice at $t = 2.4$ s; (d) horizontal slice at $t = 3.2$ s.

Figure 3: Vertical cross-sections and horizontal slices through the skeletonized LFE volume: (a) Vertical cross-section at $y = 1.3$ km; (b) vertical cross-section at $y = 6$ km; (c) horizontal slice at $t = 2.4$ s; (d) horizontal slice at $t = 3.2$ s.

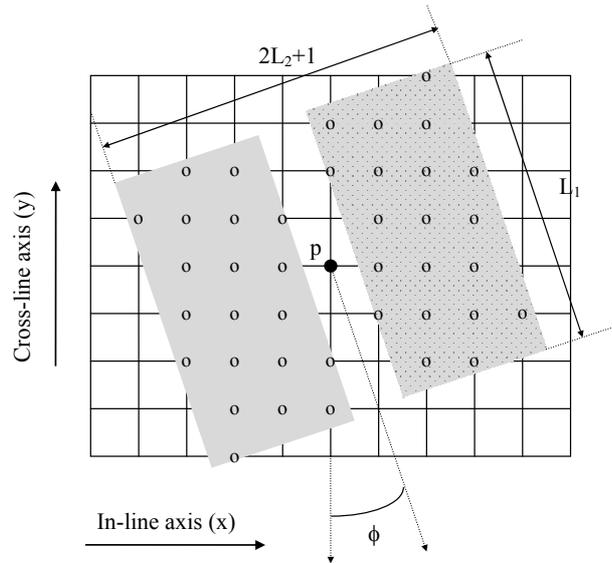
Figure 4: Three-dimensional visualization of the original seismic data and corresponding color-coded fault surfaces. Different colors correspond to different fault surface numbers. The fault surfaces are robustly extracted, separated and labeled.

Figure 5: Vertical cross-sections and horizontal slices through the Local Structural Entropy volume: (a) Vertical cross-section at $y = 1.3$ km; (b) vertical cross-section at $y = 6$ km; (c) horizontal slice at $t = 2.4$ s; (d) horizontal slice at $t = 3.2$ s.

Figure 6: (a) Eight-neighborhood notation used for binary image skeletonization. (b) Eight-neighborhood for a pixel that has only one neighbor marked as '1'. The neighbor marked as '1' either belongs to the four-neighborhood (left figure) or it is located on the corner of the eight-neighborhood (right figure). The locations of pixels that are considered "not too close" to the neighbor pixel marked as '1' are depicted by 'x'.



(a)



(b)

Figure 1: (a) A vertical cross-section and (b) horizontal slice illustrating the geometrical distribution of traces and samples used in the analysis cube. The analysis cube, consisting of two subvolumes, is centered about an analysis point $p = (t, x, y)$, and defined by length of major axis L_1 , length of minor axis $2L_2 + 1$, time duration N samples, azimuth ϕ around the in-line axis, and tilt θ from the vertical axis.

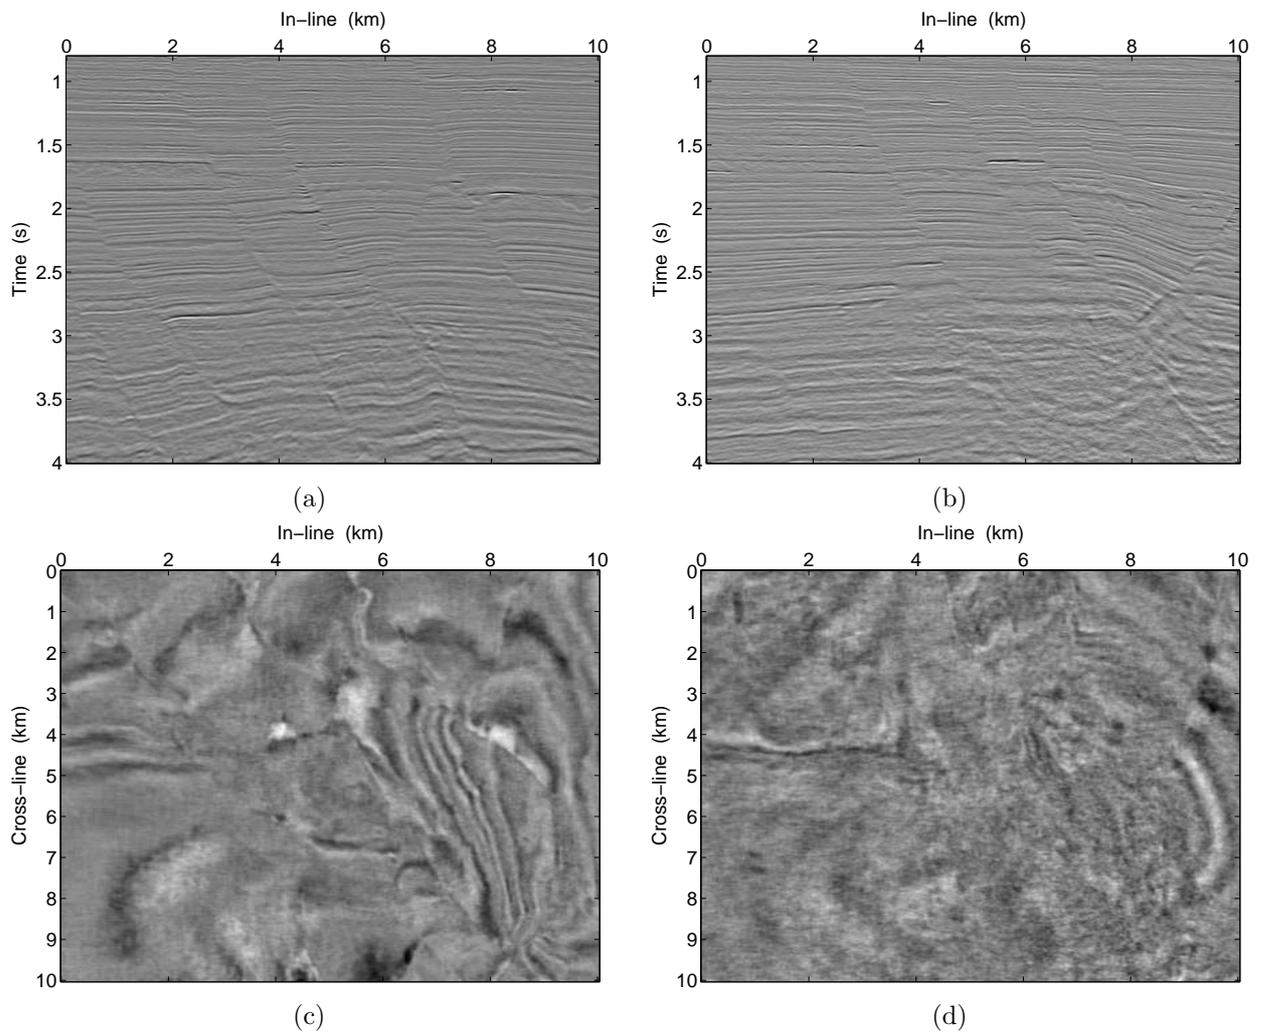


Figure 2: Vertical cross-sections and horizontal slices through the seismic data: (a) Vertical cross-section at $y = 1.3$ km; (b) vertical cross-section at $y = 8$ km; (c) horizontal slice at $t = 2.4$ s; (d) horizontal slice at $t = 3.2$ s.

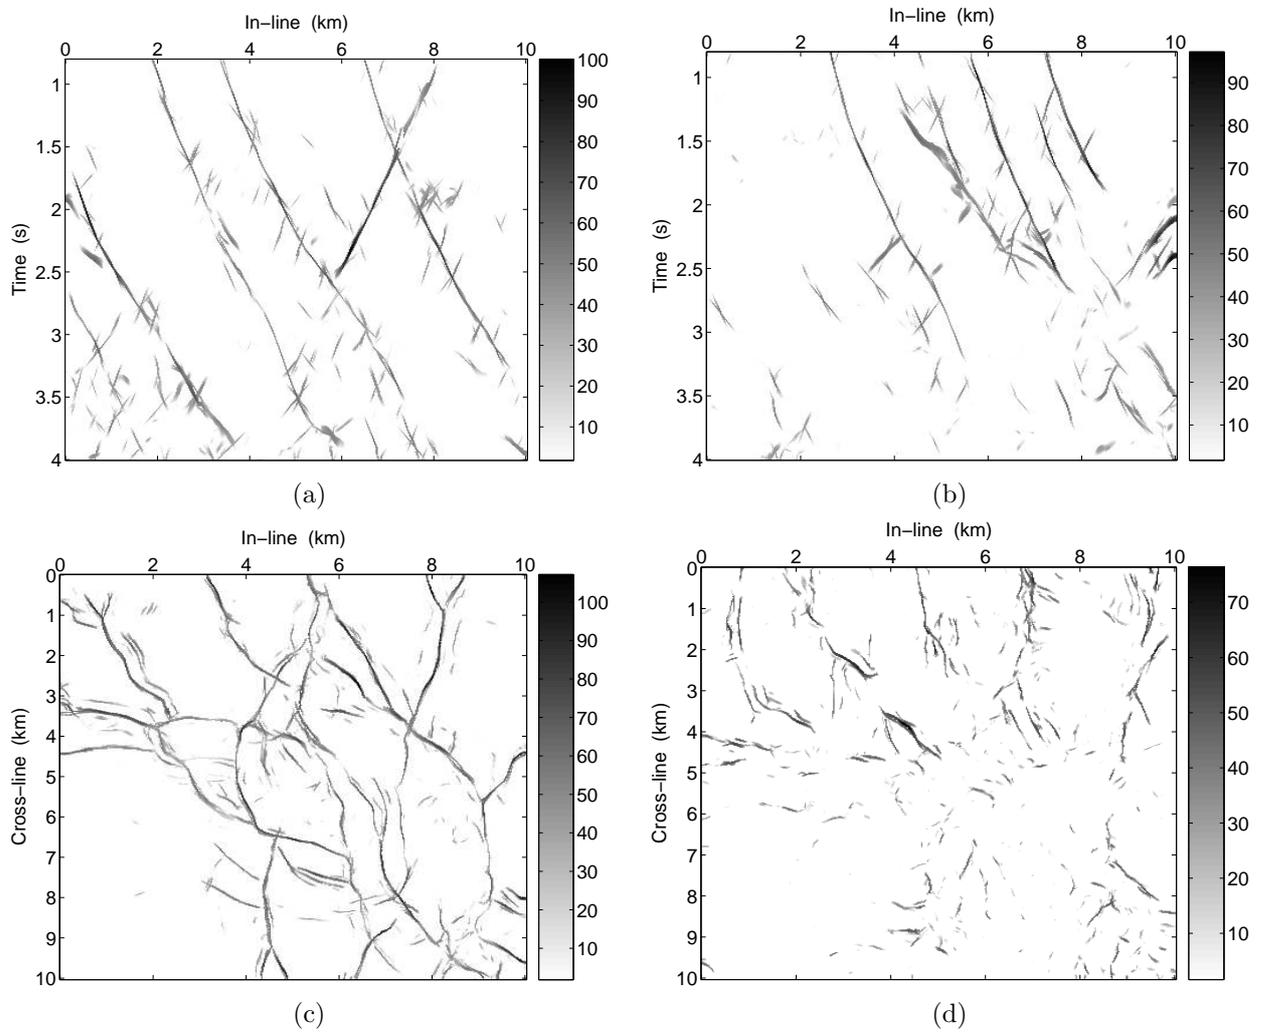


Figure 3: Vertical cross-sections and horizontal slices through the skeletonized LFE volume: (a) Vertical cross-section at $y = 1.3$ km; (b) vertical cross-section at $y = 6$ km; (c) horizontal slice at $t = 2.4$ s; (d) horizontal slice at $t = 3.2$ s.

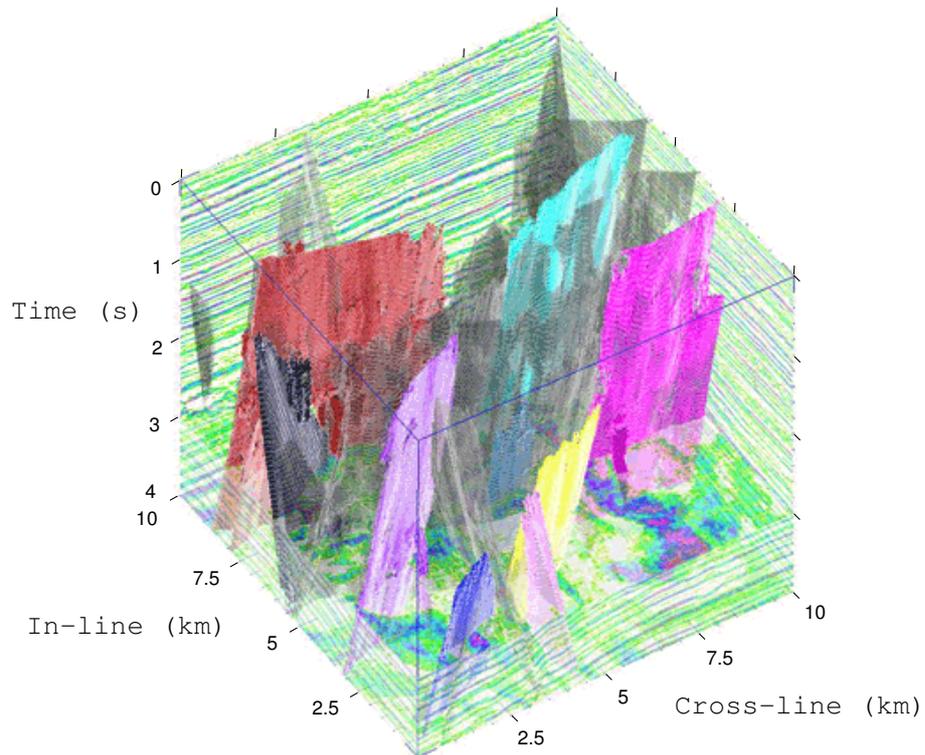


Figure 4: Three-dimensional visualization of the original seismic data and corresponding color-coded fault surfaces. Different colors correspond to different fault surface numbers. The fault surfaces are robustly extracted, separated and labeled.

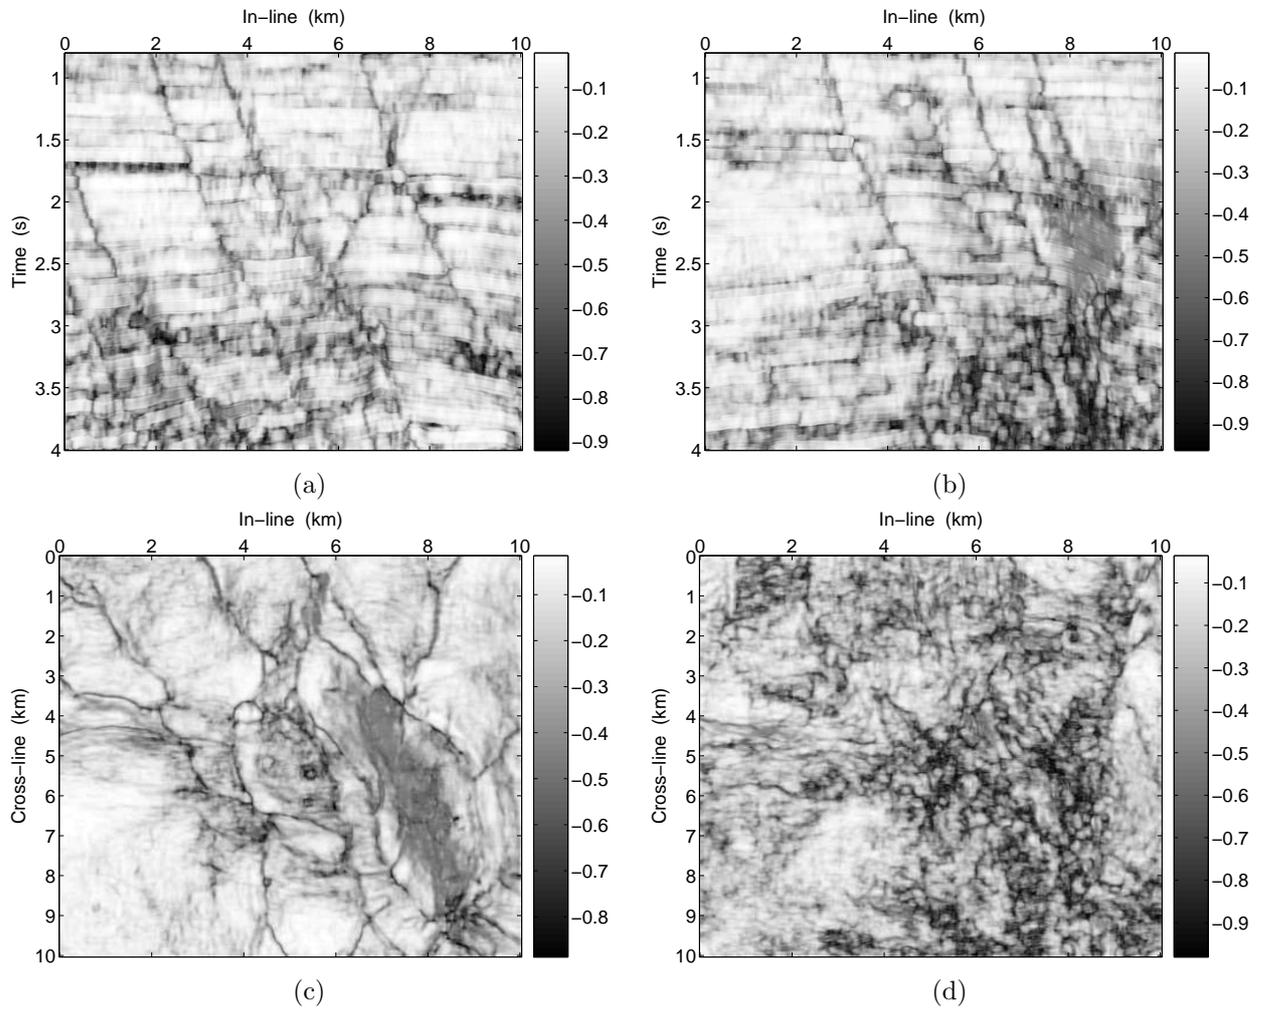


Figure 5: Vertical cross-sections and horizontal slices through the Local Structural Entropy volume: (a) Vertical cross-section at $y = 1.3$ km; (b) vertical cross-section at $y = 6$ km; (c) horizontal slice at $t = 2.4$ s; (d) horizontal slice at $t = 3.2$ s.

p_2	p_3	p_4
p_9	p_1	p_5
p_8	p_7	p_6

(a)

	Neighbor Marked as '1'	
	Center Pixel	
X	X	X

Neighbor Marked as '1'		
	Center Pixel	X
	X	X

(b)

Figure 6: (a) Eight-neighborhood notation used for binary image skeletonization. (b) Eight-neighborhood for a pixel that has only one neighbor marked as '1'. The neighbor marked as '1' either belongs to the four-neighborhood (left figure) or it is located on the corner of the eight-neighborhood (right figure). The locations of pixels that are considered "not too close" to the neighbor pixel marked as '1' are depicted by 'x'.

Israel Cohen received the B.Sc. (*Summa Cum Laude*), M.Sc. and Ph.D. degrees in electrical engineering in 1990, 1993 and 1998, respectively, all from the Technion – Israel Institute of Technology, Haifa, Israel.



From 1990 to 1998, he was a Research Scientist at RAFAEL research laboratories, Haifa, Israel Ministry of Defense. From 1998 to 2001, he was a Postdoctoral Research Associate at the Computer Science Department, Yale University, New Haven, CT. Since 2001, he has been a Senior Lecturer with the Electrical Engineering department, Technion, Israel. His research interests are statistical signal processing, analysis and modeling of acoustic signals, speech enhancement, noise estimation, microphone arrays, source localization, blind source separation, system identification and adaptive filtering.

Dr. Cohen serves as Associate Editor for the IEEE TRANSACTIONS ON SPEECH AND AUDIO PROCESSING and IEEE SIGNAL PROCESSING LETTERS, and as guest editor for a special issue of the EURASIP Journal on Applied Signal Processing on Advances in Multimicrophone Speech Processing and a special issue of the Elsevier Speech Communication Journal on Speech Enhancement.

Nicholas Coult received a B.A. in Mathematics (Magna Cum Laude, Phi Beta Kappa) from Carleton College, Northfield, Minnesota, in 1993. He received a National Science Foundation Graduate Research Fellowship to study Applied Mathematics at the University of Colorado at Boulder, and received his M.S. and Ph.D. degrees from that institution in 1995 and 1997, respectively. From 1998 to 2000 he was an Industrial Postdoctoral Associate the Institute for Mathematics and Its Applications at the University of Minnesota. From 2000 to 2005 he was an Assistant Professor of Mathematics at Augsburg College in Minneapolis, Minnesota. Since then he has been a scientific consultant for GeoEnergy, Inc., and Fast Mathematical Algorithms and Hardware Corporation. His research interests are in image processing, numerical methods for wave propagation, and high performance computing.





Anthony Vassiliou received the B.Sc. in Geodesy from the National Technical University of Athens, Greece in 1981. He received the M.Sc. and Ph.D in Civil Engineering from the University of Calgary, Canada in 1984 and 1986 respectively. From 1987 till 1992 he was a Research Geophysicist with Mobil Research and Development Corporation. From 1993 until spring of 1998 he was a Sr. Research Geophysicist with Amoco Production. In early 1998 he founded GeoEnergy, Inc. of which he is the CEO from spring 1998 till now. His research interests are seismic inversion based reservoir characterization, local time-frequency analysis and denoising, wave equation propagation.