CWP-702
May 2011



An image-guided method for automatically picking common-image-point gather locations from seismic images

Thomas A. Cullison

— Master of Science Thesis —
Mathematical & Computer Sciences

Defended on April 1, 2011

Committee Chair:	Prof. Andrzej Szymczak
Advisor:	Prof. Mahadevan Ganesh
Co-advisor:	Prof. Paul Sava
Committee members:	Prof. Dave Hale
	Dr. Scott Morton

Center for Wave Phenomena
Colorado School of Mines
Golden, Colorado 80401
(1) 303 273-3557

Abstract

Using common-image-point gathers (CIPs) instead of conventional image gathers for migration velocity analysis is an efficient approach for analyzing migrated images at locations that are related to geological structures. This approach reduces computational costs and may improve migration velocity analysis. However, manually picking CIP locations is tedious and time-consuming, especially for 3D images; therefore, an automated method for picking CIPs is desirable. To take advantage of the potential computational cost savings, CIPs should be constructed at relatively sparse locations throughout an image. Furthermore, to facilitate improved analysis, CIPs should be constructed along image features that correspond to geological features. In this thesis, a new method for automatically picking CIP locations from seismic images is proposed and demonstrated. This method uses local image properties including linearity, planarity, structure-oriented semblance, and the amplitude envelope to compute a priority map (PMAP) of seismic images, where higher PMAP values indicate better CIP locations. After the PMAP has been computed, CIP locations are picked sparsely throughout an image and along geologically related features using a greedy heuristic.

Table of Contents

Abstract i

Acknowledgments v

Chapter 1 Introduction 1

1.1 Ray-based Tomography with Limited Picking 3

1.2 An Automated Method for Picking CIPs 3

Chapter 2 Wavefield-based imaging 5

2.1 Common-image-point Gathers 6

Chapter 3 Identifying good CIP locations in seismic images 8

Chapter 4 Priority map 13

4.1 Structure Tensors 14

4.2 Linearity 15

4.3 Planarity 16

4.4 Structure-oriented Semblance 18

4.5 Linear Structure-oriented Semblance 22

4.6 Planar Structure-oriented Semblance 24

4.7 Envelope 27

4.8 Mask 27

4.9 Computing the Priority Map 29

Chapter 5 Greedy Heuristic for Picking CIP Locations 32

5.1 Greedy Algorithms 32

5.2 Greedy Heuristic for Picking CIP Locations 33

5.3 Exclusion Zones 34

5.4 Runtime Complexity 36

Chapter 6 Results 43

6.1 Sigsbee 2A Data 43

6.2 SEAM Data 43

6.3 Verifying the Quality of 3D Picks 45

Chapter 7 Conclusions 53

References 55

Acknowledgments

I would like to first acknowledge my friends and family for their support during my studies, and when I chose to leave my job and go back to school. I would also like to thank my advisors Mahadevan Ganesh and Paul Sava as well as my committee members Dave Hale, Andrzej Szymczak, and Scott Morton for their guidance and assistance with my research. They have all been incredibly helpful and insightful.

I have had a wonderful time interacting and collaborating with the students and faculty at the Center for Wave Phenomena. I would especially like to thank Diane Witters for her help with my writing and her never ending patience and understanding. I would also like to give a special thanks to Jeff Godwin, Francesco Perrone, and Tongning Yang for their helpful discussions and collaboration.

Chapter 1

Introduction

Wavefield-based depth migration is commonly used for creating accurate seismic images of complex subsurface geology. However, the accuracy of migrated images is strongly related to the accuracy of the velocity models that were used for migration. Therefore, generating accurate velocity models for migration is necessary for accurate imaging (Gray *et al.* , 2001).

Sava & Vasconcelos (2011) explain that there are two categories of methods that use seismic wavefields for estimating and improving the accuracy of velocity models. In one category of methods, an objective function is formulated prior to migration that minimizes the mismatch between data that are recorded in the field and data that are generated synthetically using approximate velocity models. These methods are typically referred to as tomography. In the other category of methods, an objective function is formulated after migration. This objective function measures and corrects for image features that indicate velocity model inaccuracies. Methods in this second category are typically known as migration velocity analysis (MVA), and these methods are relevant to this thesis.

In MVA, image attributes are analyzed for indications of inaccurate imaging. These attributes are typically represented by image extensions using image-gathers which are a collection of images for specific image locations that have been migrated for multiple seismic experiments. These gathers can be analyzed for velocity error by using the *semblance principle* which states that image-gathers are flat or well focused when velocity models are correct (Shen *et al.* , 2003; Stolk *et al.* , 2005). Using this principle, velocity models can be

built based on objective functions that optimize for image-gather features such as focusing or flattening (Al-Yahya, 1989; Symes, 2008).

The use of common-image-gathers (CIG) for MVA is standard, but a new type of image-gather called a common-image-point gather (CIP) has recently been proposed by Sava & Vasconcelos (2011). They discuss several advantages that CIPs have compared to (CIGs), but there are two main advantages. First, CIPs are constructed for a single image point; therefore, the image can influence where and how sparse CIPs are located, and because CIPs can be constructed at relatively sparse locations, they can be more computationally efficient than CIGs. In contrast to CIPs, to be computationally efficient, CIGs are typically constructed at every depth location for fixed surface coordinates which can result in the image being undersampled laterally and oversampled vertically. Second, CIPs can provide useful velocity analysis for image structures (reflectors) that have any dip. For example, the image can be used to constrain CIPs to locations where reflectors show strong coherency and simple structure, and the sparseness of these locations can be influenced by the structure of local image features (i.e., reflectors, diffractors, faults, etc.). However, when CIGs are constructed in the manner discussed above, they provide useful velocity analysis only for nearly horizontal reflectors (Sava & Vasconcelos, 2011).

Manually picking sparse CIP locations from complex seismic images can be a difficult and time-consuming task, especially for 3D images. For example, a user might need to scan through hundreds or even thousands of 2D image slices and meticulously select locations that are along simple, coherent reflectors while trying to maintain a minimum distance between selected locations. Therefore, if CIPs are used for MVA, it's essential that an automated method for picking CIP locations be used. Furthermore, because features and properties of the image should dictate the placement of CIPs, any automated method should use measurable image properties to guide the picking.

1.1 Ray-based Tomography with Limited Picking

A method that was described by Clapp (2001) which was originally intended for picking locations on reflectors for ray-based tomography can be used for picking CIP locations. This method uses measurable image properties such as dip coherency, amplitude, and semblance to constrain pick locations. Initially, locations are picked based on user-specified thresholds of dip coherence, amplitude, and distance from other points. Next, the picks are improved by iteratively relaxing the dip coherence and amplitude thresholds. For example, in the first pass, locations above the 90^{th} percentile in dip coherence and amplitude can be picked, and in the last pass, locations that are above the 50^{th} percentile can be picked. Finally, semblance is computed for each of the picked locations, and the picks that are associated with poor semblance are discarded.

If the method presented by Clapp (2001) is used for picking CIP locations, it has several drawbacks. First, it is important that CIPs be constructed near the center of reflectors (this is discussed later), but this method does not try to constrain picks to these locations. Second, this method doesn't allow the image to influence the sparsity of the picks. Instead, a user-defined parameter enforces a minimum sparsity. Third, if the picking process produces more picks than is desired, it isn't clear that the number of picks can be easily reduced based on the image-properties without also removing too many picks. For example, the number of picks could be reduced based on thresholds for either dip coherence, amplitude, semblance, or any combination of these image-properties. The picks that correspond to locations that are below these thresholds are discarded. However, this may require experimenting with different thresholds to keep from discarding too many picks.

1.2 An Automated Method for Picking CIPs

In this thesis, an automated method for picking CIP locations form seismic images is proposed and demonstrated. This method is guided by measurable image properties, and

it does not suffer from the drawbacks discussed above. The method has two stages. In the first stage, a priority map (PMAP) of a seismic image is made. In the second stage, a greedy picking heuristic is used to automatically pick CIP locations. The PMAP guides and constrains the picking of CIP locations, and the greedy picker provides the means to control the sparseness of the picks. The greedy picker also maintains an ordering of the picks based on their corresponding PMAP values. The PMAP contains a value for each point of a seismic image where higher values indicate better CIP locations, and it is constructed by combining measurable image properties. These properties are derived from structure tensor fields, structure-oriented semblance, and the amplitude envelope. The greedy picker is guided by either the PMAP or both the PMAP and the structure of local image features (i.e., reflector structure).

The following chapter gives a brief overview of the imaging conditions used in wavefield-based migration methods and of how CIPs can be used for MVA. Chapter 3 discusses the image features and properties that indicate good CIP locations. Chapter 4 discusses the PMAP and how it is constructed from measurable image properties. In Chapter 5, the greedy picking heuristic is described, and in Chapter 6, examples of CIP picks that were obtained using the automated picking method on 2D and 3D benchmark images are shown and validated.

Chapter 2

Wavefield-based imaging

For wavefield-based depth migration methods, it is standard to assume the single scattering approximation which states that waves in the subsurface scatter only once at discontinuities which are subsurface locations where there is a change in acoustic impedance (i.e., the contacts between geologic units or structures). Under this approximation, subsurface wavefields are typically split into source and receiver wavefields. The source wavefields originate from seismic sources and propagate through the subsurface without interacting (scattering) with discontinuities. The receiver wavefields are assumed to originate at discontinuities and propagate to receivers without interacting with other discontinuities within the subsurface (Berkhout & Pao, 1982; Claerbout, 1985). It is also standard to assume that when the velocity model used for migration is accurate, the source and receiver wavefields will kinematically match at discontinuities, and that any mismatch is due to an inaccurate velocity model (Sava & Vasconcelos, 2011).

In this context, the imaging process can be broken into two steps: the wavefield reconstruction and the imaging condition. In the wavefield reconstruction step, the source and receiver wavefields W_s and W_r for a single seismic experiment (shot) are reconstructed by a numerical solution to the wave equation for which data represent both the boundary and initial conditions. These wavefields can be represented as functions of space and time.

In the imaging condition step, seismic images are constructed by evaluating the kinematic match between the source and receiver wavefields from multiple shots and then sum-

ming the result. The kinematic match between source and receiver wavefields is evaluated using an imaging condition that operates along the space and time axes. The conventional imaging condition (cIC) which is used to extract a conventional image can be represented by a zero cross-correlation lag between the source and receiver wavefields W_s and W_r (Claerbout, 1985). This type of cIC is defined as

$$R\left(\mathbf{x}\right) = \sum_{shots} \sum_{t} W_s\left(\mathbf{x}, t\right) W_r\left(\mathbf{x}, t\right) \ , \tag{2.1}$$

where, R is the migrated image. Other imaging conditions such as deconvolution between W_s and W_r can be used as well, but a discussion of these conditions is not relevant to this thesis.

An extended imaging condition (eIC) is used to extract extended images such as CIPs and CIGs, and it is typically computed using a non-zero cross-correlation lag between the reconstructed wavefields W_s and W_r (Sava & Vasconcelos, 2011). The eIC is defined as

$$R\left(\mathbf{x}, \boldsymbol{\lambda}, \tau\right) = \sum_{shots} \sum_{t} W_s\left(\mathbf{x} - \boldsymbol{\lambda}, t - \tau\right) W_r\left(\mathbf{x} - \boldsymbol{\lambda}, t - \tau\right) \ , \tag{2.2}$$

where $\boldsymbol{\lambda} = \{\lambda_x, \lambda_y, \lambda_z\}$ represents the space lags and τ represents the time lag. Notice that (2.1) is a special case of (2.2) when $\boldsymbol{\lambda} = \mathbf{0}$ and $\tau = 0$. The extended image defined by (2.2) can be analyzed for velocity inaccuracies that can be used to improve the velocity model (Sava & Vasconcelos, 2011). This analysis is based on the assumption that errors in the reconstructed wavefields are the result of inaccurate velocity. Under this assumption, extended images can be used to correct the velocity based on the semblance properties (i.e., focusing) which are characterized by the space and time lags.

2.1 Common-image-point Gathers

Common-image-point gathers (CIPs) are extended images that are constructed at fixed positions of space, and they can be used for MVA to analyze and correct for velocity model

inaccuracies (Sava & Vasconcelos, 2011). However, accurate velocity analysis using CIPs may only be applicable when several assumptions are valid. It is assumed that source and receiver wavefields as well as the reflector (the area being imaged) can be approximated by planes within the vicinity of the CIP location. Furthermore, it is assumed that the velocity is locally constant in this vicinity as well. When these assumptions are not valid for a particular CIP location, then velocity analysis using this CIP can be challenging. To the author's knowledge, a robust understanding of the consequences of using CIPs for velocity analysis when the above assumptions are not valid is an open area of research.

A method for using CIPs for MVA is described by Yang & Sava (2010). This method uses differential semblance optimization (DSO) to formulate objective functions that minimize velocity errors based on the CIP characteristics that indicate accurate velocity. In DSO, the semblance principle is exploited for velocity analysis. Yang & Sava (2010) explain that focusing characteristics within CIPs can indicate accurate velocity, and that these characteristics are related to the semblance principle. The full theory behind these characteristics is beyond the scope of this thesis, but a visual understanding is presented and discussed in the following chapter.

Chapter 3

Identifying good CIP locations in seismic images

To understand where CIPs should be constructed within a seismic image, it helps to look at some examples of CIPs and to understand some of the assumptions behind using CIPs for MVA. Figure 3.1 shows an image of the Sigsbee-2A data (from the Subsalt Multiples Attenuation and Reduction Technology Joint Venture) which were migrated using the correct velocity. This image shows geologic units which are represented by the numerous curvilinear features, a large salt body which starts at about 10 km in x and continues to 26 km, and a salt water body which is the smooth structureless feature at the upper left. The locations of the CIPs shown in Figures 3.2–3.5 are plotted in Figure 3.1. The CIPs at points A and B in Figure 3.1 are shown in Figure 3.2 and 3.3, respectively. The locations of these two CIPs share a similar quality that is of importance for MVA. Notice that both CIPs are located near the center of simple, coherent reflectors, and that both CIPs have strong focusing near the zero-crossing of the λ_x and τ axes. This focusing property is useful for MVA (Sava & Vasconcelos, 2011).

As previously discussed, when CIPs are used to analyze velocity errors, it is assumed that the wavefields and reflectors are locally planar in the immediate regions that surround the CIP locations. These assumptions imply that there may be image locations where the construction of CIPs will not be useful for MVA. This includes locations where the subsurface may have strong velocity discontinuities such as at the interface of a salt body, at diffractors, or along faults because of discontinuous geologic units. In these locations

the locally planar assumption breaks down if the velocity discontinuities are strong enough. Therefore, picking at these locations should be avoided. An example of a CIP constructed at location where the planar assumption breaks down is shown in Figure 3.4. This CIP was constructed at point C in Figure 3.1 which is located at the edge of the salt body. Notice the lack of focusing near the zero-crossing of the λ_x and $\tau(s)$ axes.

When a CIP is constructed off the center of a reflector, focusing may also not occur near the zero-crossing of the λ_x and τ axes even if the wavefields have been reconstructed using the correct velocity. Therefore, CIPs should be constructed near the center of reflectors. This lack of focusing is evident when comparing the CIP constructed at point A (CIP-A) with the CIP shown in Figure 3.5 which was constructed a little ways below point A. Because the CIP in Figure 3.5 was constructed below the reflector, there isn't focusing near the zero-crossing. Instead, the focused feature in CIP-A shows up as a hyperbolic-like feature in Figure 3.5. Notice that the feature has shifted slightly above the zero crossing. This behavior is also shown in Figure 3.2 by the hyperbolic-like features that are above and below $\tau = 0$ and centered around $\lambda_x = 0$. These features are related to the reflectors that are above and below the location of CIP-A.

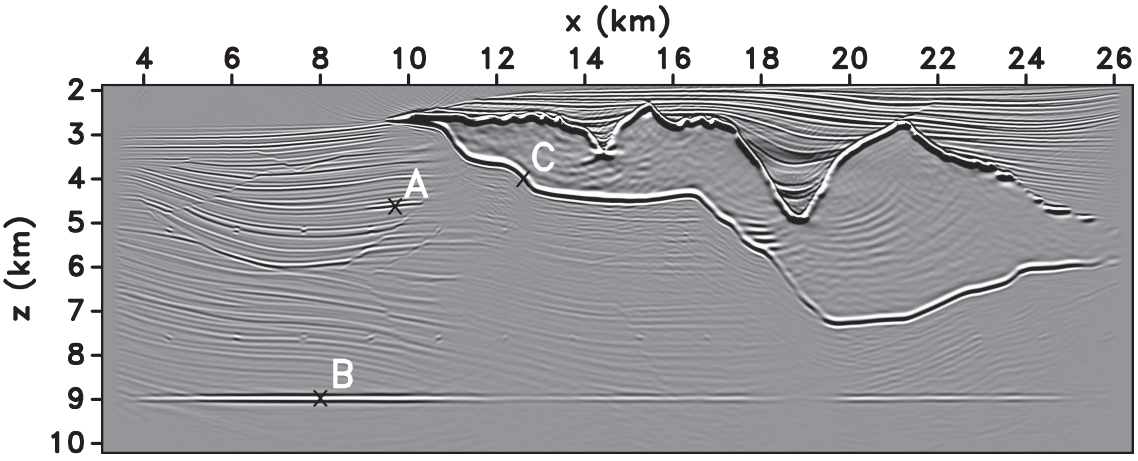


Figure 3.1. An acoustically migrated image of the Sigsbee-2A model. CIP locations are marked with an X.

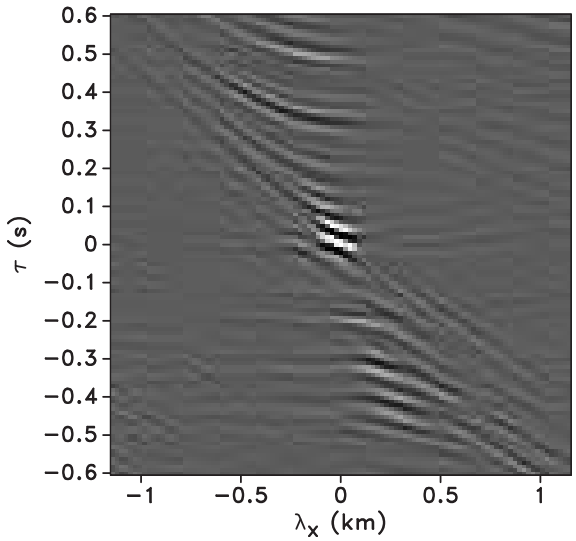


Figure 3.2. This CIP was constructed at point A in Figure 3.1. This point is located at the center of a coherent reflector with simple structure.

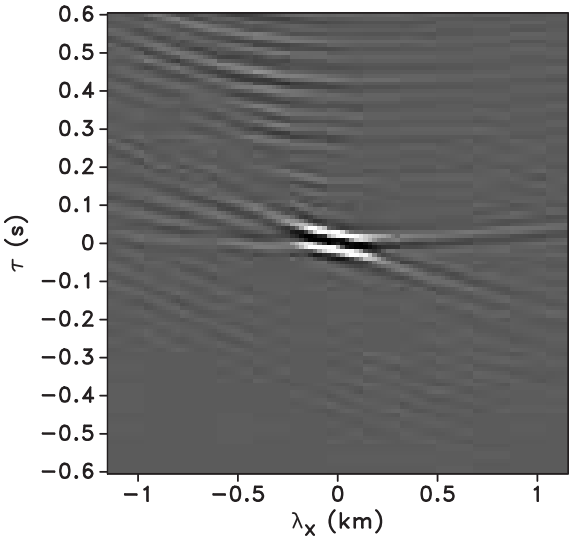


Figure 3.3. This CIP was constructed at point B in Figure 3.1. This point is located at the center of a coherent reflector with simple structure.

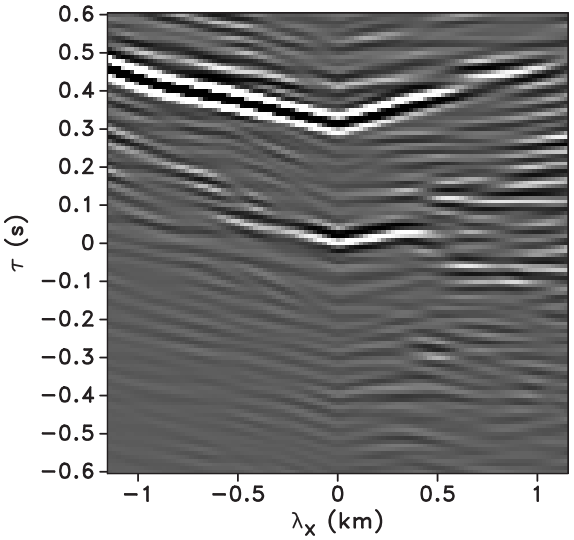


Figure 3.4. This CIP was constructed at point C in Figure 3.1. This point is located at the edge of a salt body where there is a strong discontinuity in the subsurface velocity.

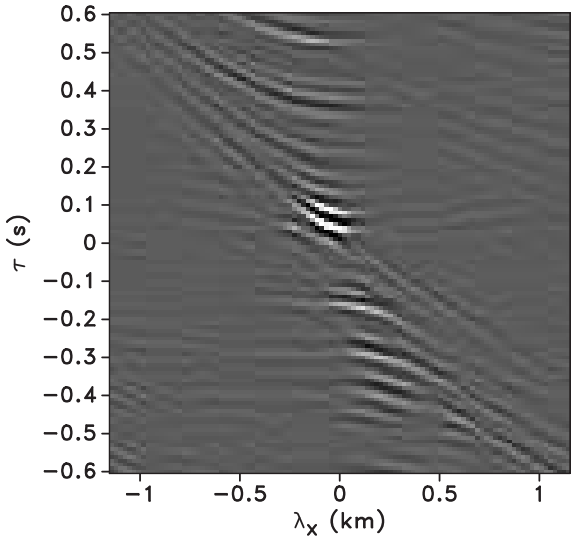


Figure 3.5. This CIP was constructed just below point A in Figure 3.1 and is located just below the center of a reflector.

Chapter 4

Priority map

As previously discussed, the construction of CIPs at locations where there are strong velocity discontinuities should be avoided, and the construction of CIPs near the center of reflectors should be favored. Therefore, determining whether or not a location in a seismic image (a pixel in 2D images or a voxel in 3D images) is near the center of a coherent reflector that has simple structure is an important problem. This problem is addressed in this thesis by first creating a priority map (PMAP) of the image which emphasizes locations that have the above qualities, and then using a greedy picking heuristic which picks CIP locations by favoring the locations that have more of these qualities.

The PMAP assigns a value to each image location, where greater PMAP values indicate locations that have more of the above qualities. This assignment is important because the greedy picking heuristic will favor locations that have greater PMAP values, and completely ignore locations that have PMAP values equal to zero. The PMAP is constructed by multiplying together several measurable image properties including linearity (for 2D images), planarity (for 3D images), structure-oriented semblance, and the amplitude envelope. These image properties act as proxies for the desired qualities discussed above, and by multiplying these properties together, greater emphasis is placed on the locations that have the best combinations of these image properties. Furthermore, a manually defined mask can also be applied to the PMAP to exclude the construction of CIPs at locations that are within specific sections of an image such as within salt or water bodies. The linearity and

planarity of 2D and 3D seismic images, respectively, are computed using structure tensors. These image properties provide a measure of the structural complexity of reflectors in an image, where linearity emphasizes locations that are within curvilinear reflectors of a 2D image, and the planarity emphasizes locations that are within curvy-planar reflectors of a 3D image. Structure-oriented semblance, which emphasizes image locations that are within coherent reflectors, is also computed using structure tensors. The amplitude envelope of a seismic image provides a measure of a location's nearness to the center of a reflector.

4.1 Structure Tensors

To get a measurement of the linearity of a 2D image or the planarity of a 3D image, a structure tensor field for the image is computed. Structure tensors provide a means for obtaining local orientations of anisotropic image features (Vliet & Verbeek, 1995). A structure tensor is a smoothed outerproduct of image gradients, and it is often referred to as the *gradient squared tensor* (Fehmers & Hocker, 2003; Hale, 2009b). The structure tensor field $\mathbf{T}(\mathbf{x})$ of an image is defined as

$$\mathbf{T}(\mathbf{x}) = \langle \mathbf{g}(\mathbf{x})\mathbf{g}^T(\mathbf{x}) \rangle \ , \tag{4.1}$$

where $\mathbf{g}(\mathbf{x})$ is the gradient of the image and $\langle \cdot \rangle$ denotes Gaussian smoothing in all coordinate directions. The gradient contains estimates of both the magnitude and the direction of the greatest change of the values (pixel or voxel brightness) of the image. Structure tensors contain similar estimates of magnitude and direction; however, due to the Gaussian smoothing, the magnitudes and directional vectors that can be derived from a structure tensor are weighted averages.

To get local measurements of linearity and planarity from structure tensors, the eigen decomposition of each tensor is computed (Hale, 2009b). The eigen decomposition of a 2D structure tensor $\mathbf{T}$ is

$$\mathbf{T} = \lambda_u \mathbf{u}\mathbf{u}^T + \lambda_v \mathbf{v}\mathbf{v}^T \ , \tag{4.2}$$

where λ_u and λ_v are non-negative eigenvalues of $\mathbf{T}$ and $\mathbf{u}$ and $\mathbf{v}$ are the respective eigenvectors. The eigenvalues of $\mathbf{T}$ are sorted and labeled such that

$$\lambda_u \geq \lambda_v \geq 0 \ . \tag{4.3}$$

Under this sorting convention, the eigenvectors $\mathbf{u}$ are aligned in the direction where the image gradient is greatest. Therefore, eigenvectors $\mathbf{u}$ are orthogonal to local linear image features and eigenvectors $\mathbf{v}$ are parallel to these features (Hale, 2009b).

The eigen decomposition of a 3D structure tensor $\mathbf{T}$ is

$$\mathbf{T} = \lambda_u \mathbf{u}\mathbf{u}^T + \lambda_v \mathbf{v}\mathbf{v}^T + \lambda_w \mathbf{w}\mathbf{w}^T \ , \tag{4.4}$$

where $\mathbf{u}$, $\mathbf{v}$, and $\mathbf{w}$ are the eigenvectors that correspond to the non-negative eigenvalues λ_u, λ_v, and λ_w of $\mathbf{T}$. Again, the eigenvalues of $\mathbf{T}$ are sorted and labeled such that

$$\lambda_u \geq \lambda_v \geq \lambda_w \geq 0 \ . \tag{4.5}$$

Similar to the 2D case, the eigenvectors $\mathbf{u}$ are aligned in the direction where the image gradient is greatest, and the eigenvectors $\mathbf{w}$ are aligned in the direction where the image gradient is weakest. Therefore, eigenvectors $\mathbf{u}$ are orthogonal to local planar image features, eigenvectors $\mathbf{w}$ are parallel to linear image features, and both eigenvectors $\mathbf{v}$ and $\mathbf{w}$ are parallel to planar features (Hale, 2009b).

4.2 Linearity

The eigenvalues λ_u and λ_v that are derived from a structure tensor provide a measure of linearity of local image features. As described by Hale (2009b), the linearity measure λ_l

is defined as

$$\lambda_l = \frac{\lambda_u - \lambda_v}{\lambda_u} \ , \tag{4.6}$$

and hence using (4.3), $0 \leq \lambda_l \leq 1$.

For convenience, the image of the Sigsbee-2A data is shown again in Figure 4.1. Figure 4.2 shows a visualization of the linearity computed for the Sigsbee image shown in Figure 4.1. The brighter areas in Figure 4.2 correspond to regions in the Sigsbee image that have curvilinear features. The linearity measure is used to constrain CIP locations to these regions (i.e., regions that have simple linear structure). Notice that the linearity values near complex structures such as faults and diffractors are typically small (fairly dim). However, these values are sometimes non-zero. Therefore, it is possible that the greedy picking heuristic will pick locations in these areas. This behavior is undesirable, but it can be mitigated by clipping the linearity values below some threshold to zero. This clipping better constrains the placements of CIPs to regions in the image that have greater linearity. Figure 4.3 shows a visualization of the linearity after clipping values below 0.9 to zero. The black regions in Figure 4.3 correspond to values that have been clipped to zero, and CIPs cannot be constructed in these locations. This clipping threshold is a user-defined parameter which is based on the visual observation of its result.

4.3 Planarity

The eigenvalues λ_u, λ_v, and λ_w that are derived for each structure tensor provide a measure of planarity of local image features. As described by Hale (2009b), the planarity measure λ_p is defined as

$$\lambda_p = \frac{\lambda_u - \lambda_v}{\lambda_u} \ , \tag{4.7}$$

where $0 \leq \lambda_p \leq 1$. Figure 4.4 shows a migrated image of the SEG Advanced Model-

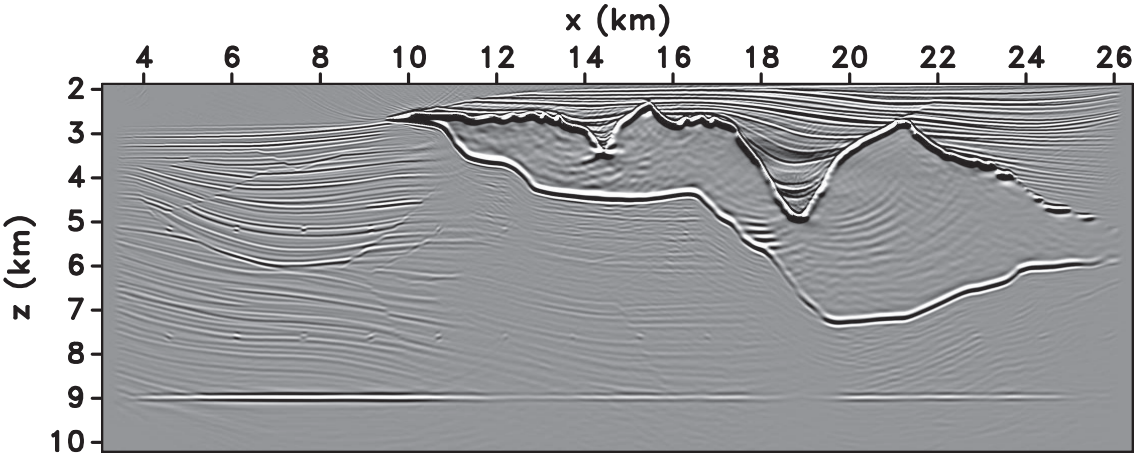


Figure 4.1. A migrated image of the Sigsbee-2A data.

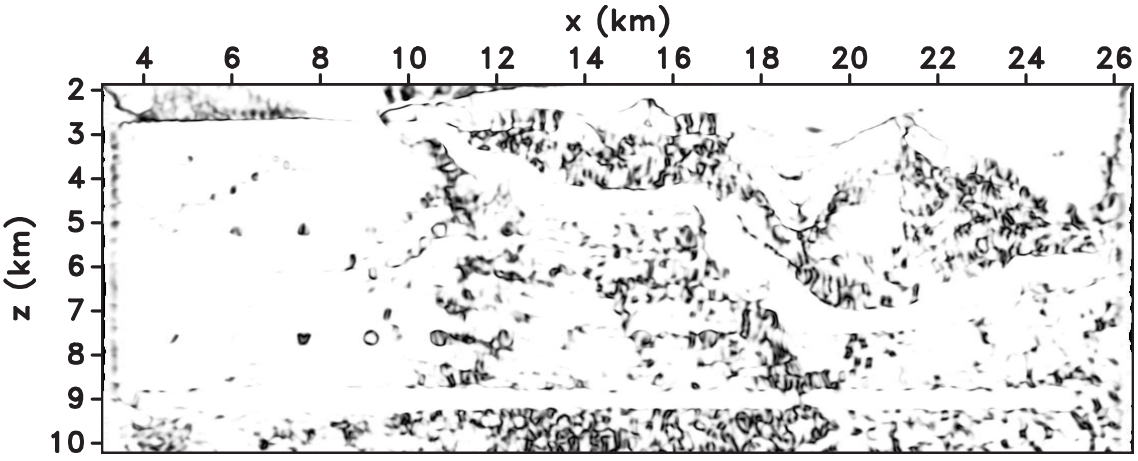


Figure 4.2. A visualization of the linearity of the Sigsbee image. Brighter values correspond to greater linearity.

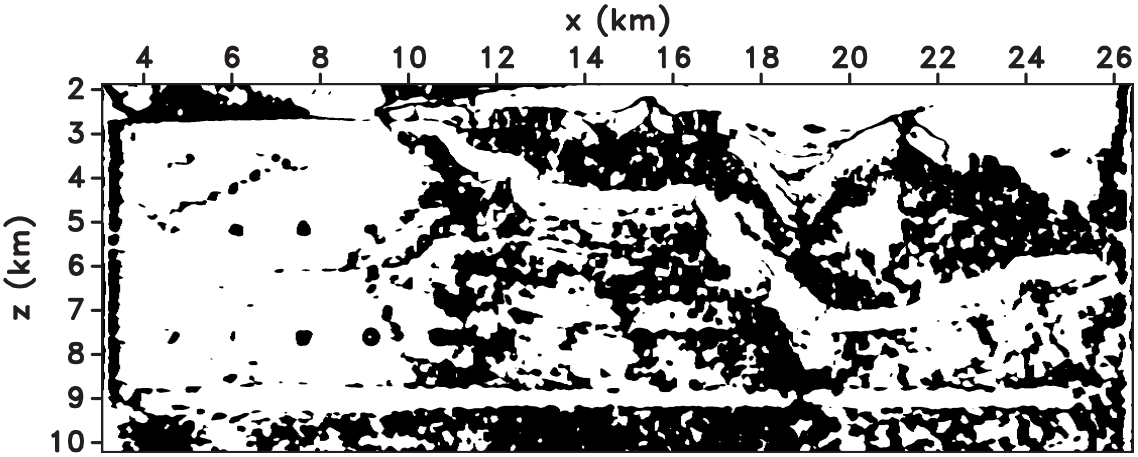


Figure 4.3. A visualization of the linearity after clipping. The clipping increases the extent of the regions where CIPs cannot be picked (the regions where the linearity values are zero).

ing (SEAM) data, and Figure 4.5 shows a visualization of the planarity computed for the SEAM image. For 3D images, the planarity measure can be used to constrain CIP locations to regions that have greater planarity. The regions in Figure 4.5 that have simple planar reflectors correspond to greater planarity values (the greater the planarity the brighter the value), and the locations near complex structures typically correspond to lower values. However, because values that are near complex structures are sometimes non-zero, CIPs locations can be picked in these areas which is undesirable. This behavior can be mitigated by clipping the planarity values below a user defined threshold to zero, which better constrains the placements of CIPs to regions that have greater planarity. Figure 4.6 shows a visualization of the planarity after clipping the values below 0.75 to zero. As with the linearity clipping, the black regions in Figure 4.6 correspond to values that have been clipped to zero, and CIPs cannot be constructed in these locations. This clipping parameter is also a user-defined parameter based on the visual observation of its result.

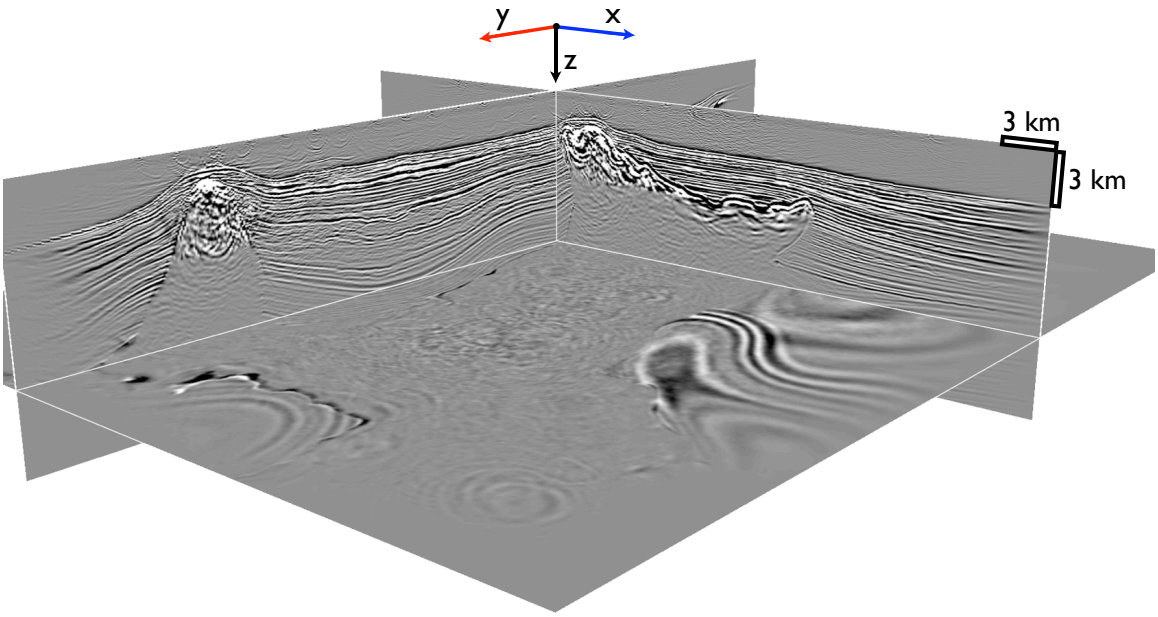


Figure 4.4. A migrated image of the SEAM model.

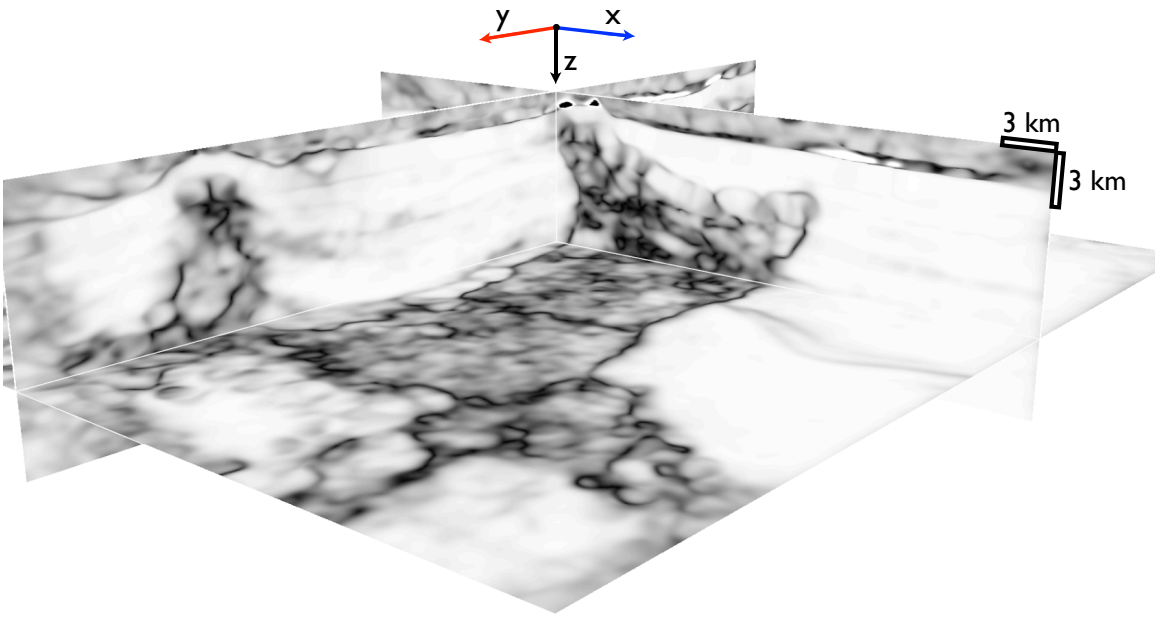


Figure 4.5. A visualization of the planarity of the SEAM image. Brighter values correspond to greater planarity.

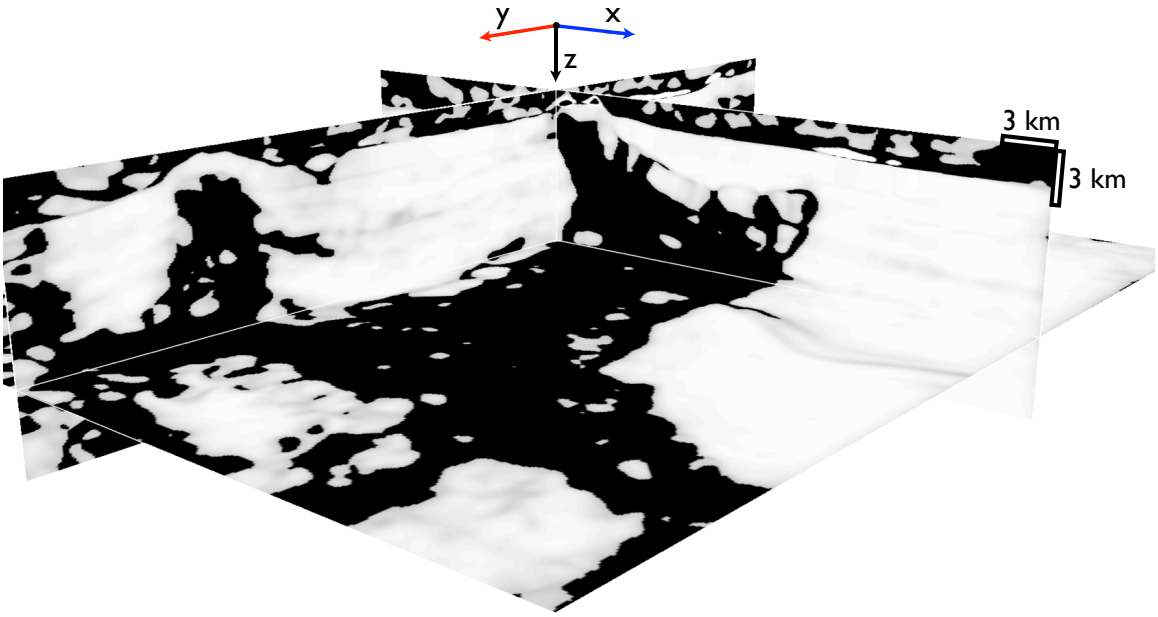


Figure 4.6. A visualization of the planarity after clipping. The clipping increases the extent of regions where CIPs cannot be picked (the black regions where the planarity values are zero).

4.4 Structure-oriented Semblance

Structure-oriented semblance provides a normalized measure of the structurally-dependent coherence of an image, and it can be used to constrain CIP locations to image features (reflectors) that are coherent and have simple structure. The method used in this thesis to compute structure-oriented semblance is described by Hale (2009b).

4.5 Linear Structure-oriented Semblance

For 2D images, linear structure-oriented semblance is used to constrain CIP locations. The linear structure-oriented semblance of a 2D image f is defined as

$$s_{v,u} = \frac{\langle \langle f \rangle_v^2 \rangle_u}{\langle \langle f^2 \rangle_v \rangle_u} \,, \tag{4.8}$$

where $\langle \cdot \rangle_v$ denotes smoothing in the local structure direction and $\langle \cdot \rangle_u$ denotes smoothing in the cross-structure direction. Both smoothing directions are determined by structure tensors, and the amount of smoothing in each direction is determined by the user. Figure 4.7 shows a visualization of the linear structure-oriented semblance computed for the Sigsbee image in Figure 4.1. Notice that the bright areas in Figure 4.7 correspond to structurally coherent reflectors, and that the darker areas correspond to faults, diffractors, and other complex structures. This semblance measure can be used to constrain CIP locations to regions in 2D images that have higher semblance values (i.e., coherent reflectors with simple structure). However, as with the linearity measure, the locations near faults, diffractors, reflector truncations, etc., can have semblance values that are non-zero, but CIP locations can be better constrained if the semblance values below some threshold are clipped to zero. Figure 4.8 shows a visualization of the linear structure-oriented semblance after clipping values that are below 0.8 to zero. The black regions correspond to values that have been clipped to zero, and CIPs cannot be constructed in these locations. This clipping threshold is another user-defined parameter that is based the visual observation of its result.

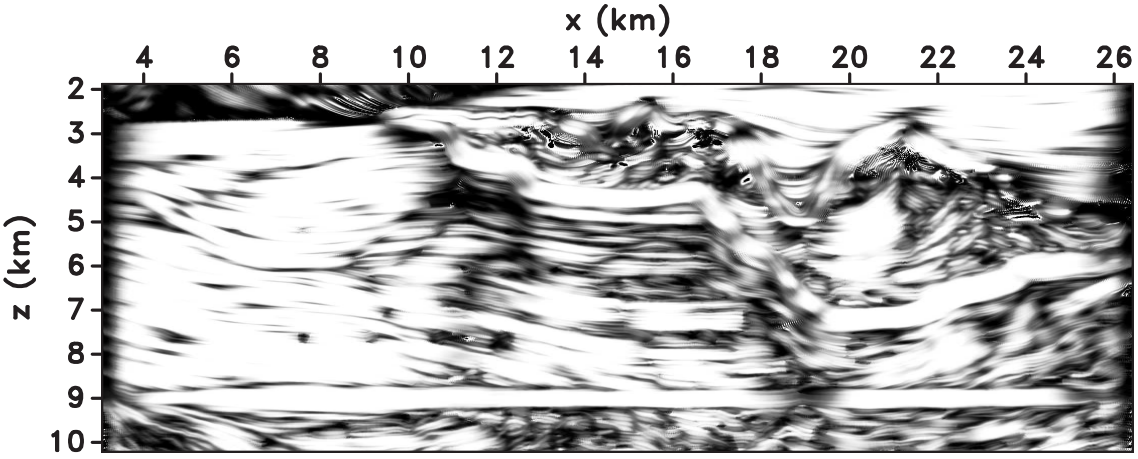


Figure 4.7. A visualization of the linear structure-oriented semblance of the Sigsbee image. Brighter values correspond to greater semblance.

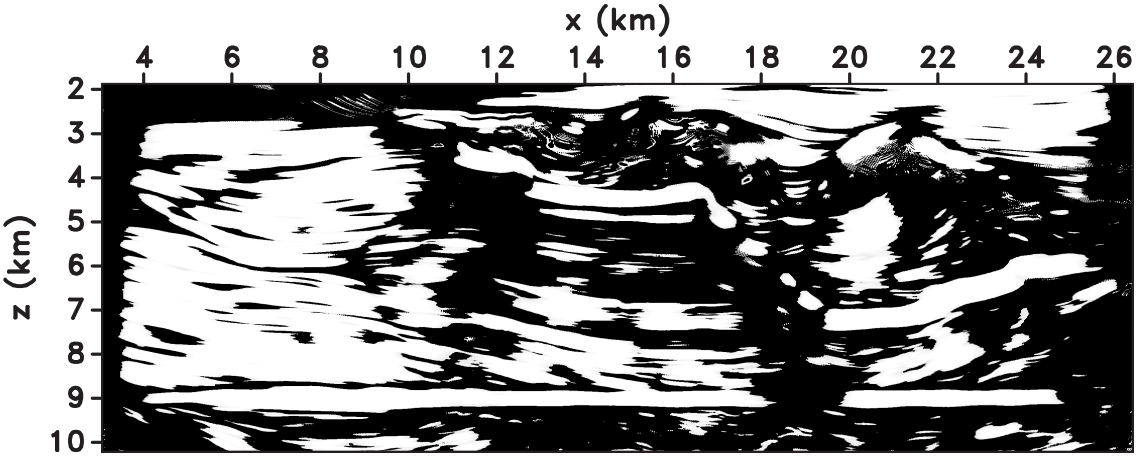


Figure 4.8. A visualization of the linear structure-oriented semblance after clipping.

4.6 Planar Structure-oriented Semblance

For 3D images, planar structure-oriented semblance is used to constrain CIP locations. The planar structure-oriented semblance of a 3D image f is defined as

$$s_{vw,u} = \frac{\langle \langle f \rangle_{vw}^2 \rangle_u}{\langle \langle f^2 \rangle_{vw} \rangle_u} \,, \tag{4.9}$$

where $\langle \cdot \rangle_{vw}$ denotes smoothing in the local structure directions (planar directions) and $\langle \cdot \rangle_u$ denotes smoothing in the cross-structure direction. Similar to linear structure-oriented semblance, the smoothing directions are determined by the structure tensors, and the amount of smoothing in each direction is determined by the user. Figure 4.9 shows a visualization of the planar structure-oriented semblance computed for the SEAM image shown in Figure 4.4. As with the linear structure-oriented semblance measure, the bright areas in Figure 4.9 correspond to structurally coherent reflectors, and the darker areas correspond to faults, diffractors, and other complex structures. This semblance measure is used to constrains CIP locations to regions that have higher semblance values (i.e., coherent reflectors with simple structure). However, locations that are near faults, diffractors, reflector truncations, etc., may have semblance values that are non-zero which is undesirable. However, clipping the semblance values below some threshold to zero can better constrain CIP locations. Figure 4.10 shows a visualization of the planar structure-oriented semblance after clipping values that are below 0.3 to zero. The black regions correspond to values that have been clipped to zero, and CIPs cannot be constructed in these regions. Again, this clipping parameter is defined based on the visual observation of its result. Automating the selection of this clipping parameter and those previously discussed, would be an interesting future work.

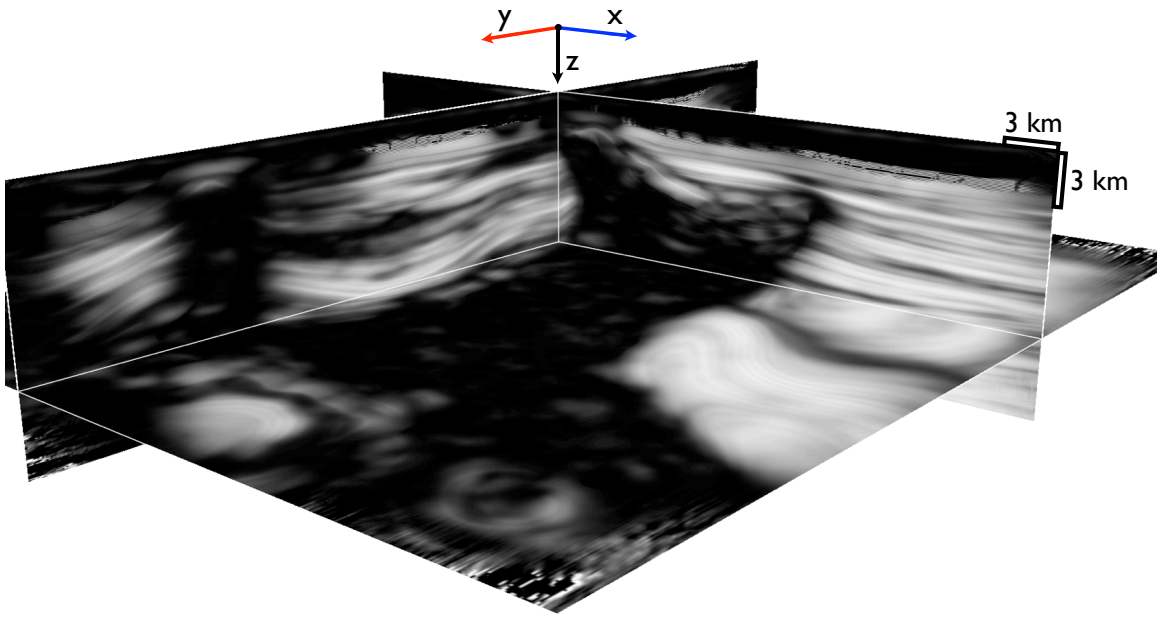


Figure 4.9. A visualization of the planar structure-oriented semblance of the SEAM image. Brighter values correspond to greater semblance.

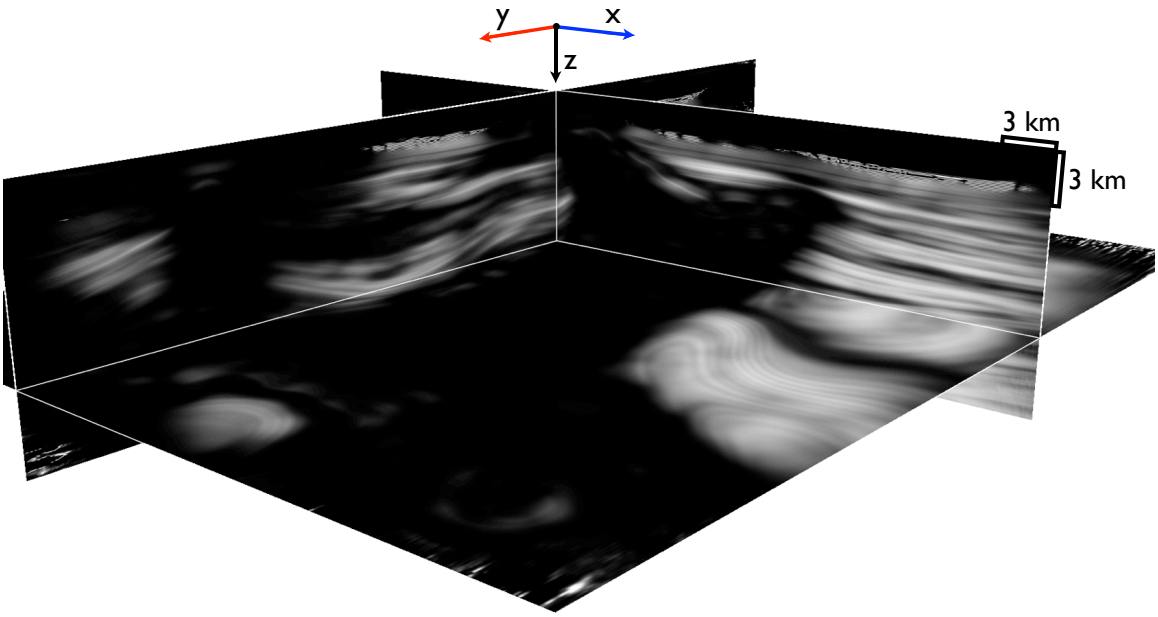


Figure 4.10. A visualization of the planar structure-oriented semblance after clipping.

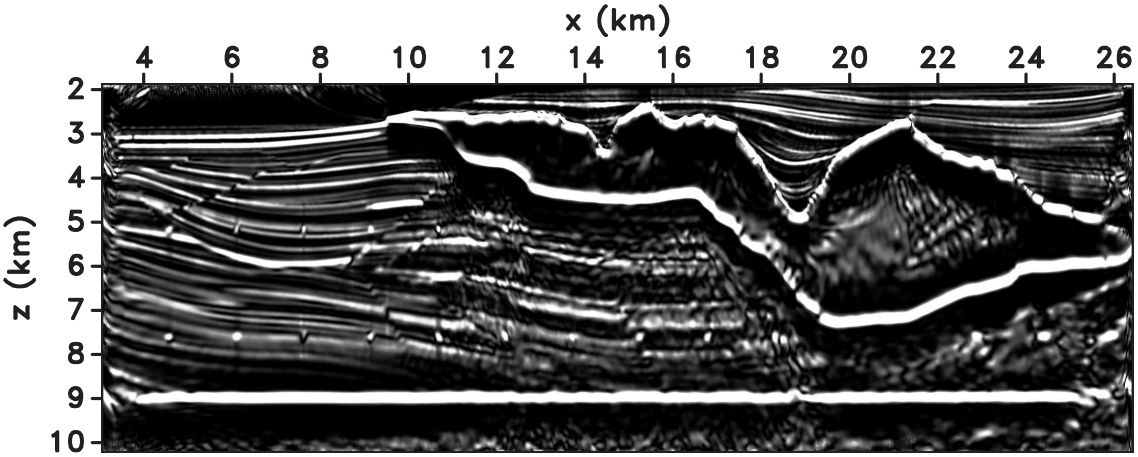


Figure 4.11. The envelope of the Sigsbee image. Brighter values are near the center of reflectors.

4.7 Envelope

To provide a measure of the distance to the center of a reflector, the vertical amplitude envelope of an image is computed, and then an automatic gain control is applied to the envelope to boost the amplitudes of weaker reflectors. If the weaker amplitudes were not boosted, CIP locations at strong reflectors would be overly favored by the greedy picking heuristic which is discussed in the following chapter. Figure 4.11 shows the envelope of the Sigsbee image, and Figure 4.12 shows the envelope of the SEAM image.

4.8 Mask

It is possible to manually exclude the construction of CIPs from specific sections of an image (e.g., water or salt bodies) by creating a mask. For example, Figure 4.13 shows a mask created for the Sigsbee image and Figure 4.14 shows a mask created for the SEAM image. Both masks were designed to manually exclude the water and salt bodies. The black regions in the mask indicate locations where CIPs cannot be constructed. The values in these regions are set to zero. The white regions are locations where the linearity or the

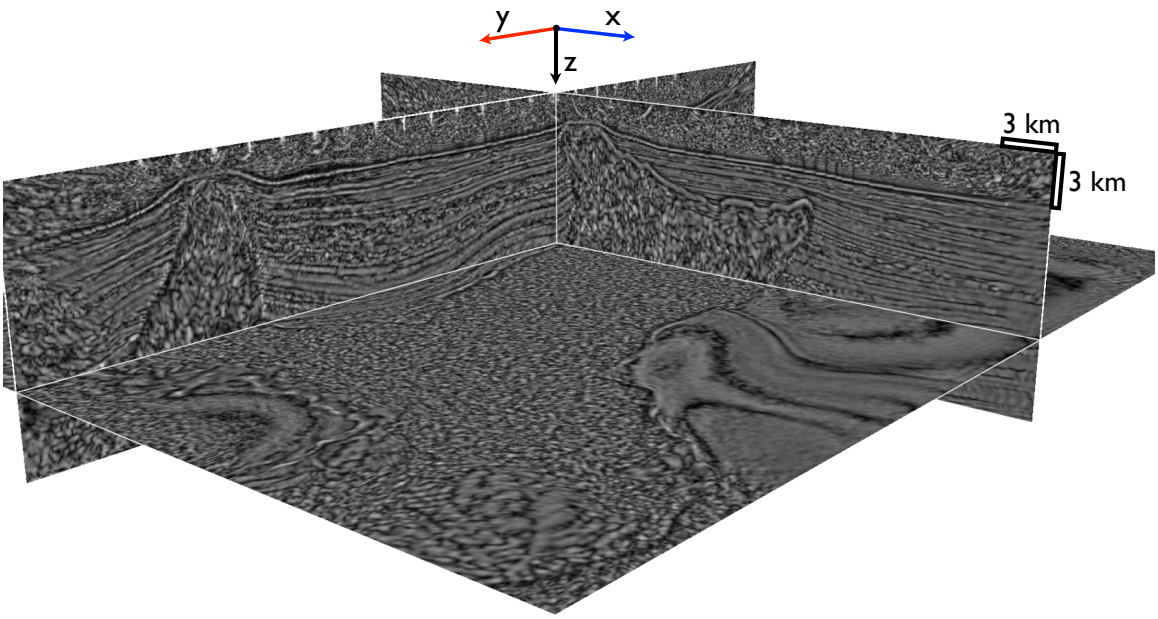


Figure 4.12. The envelope of the SEAM image. Brighter values are near the center of reflectors.

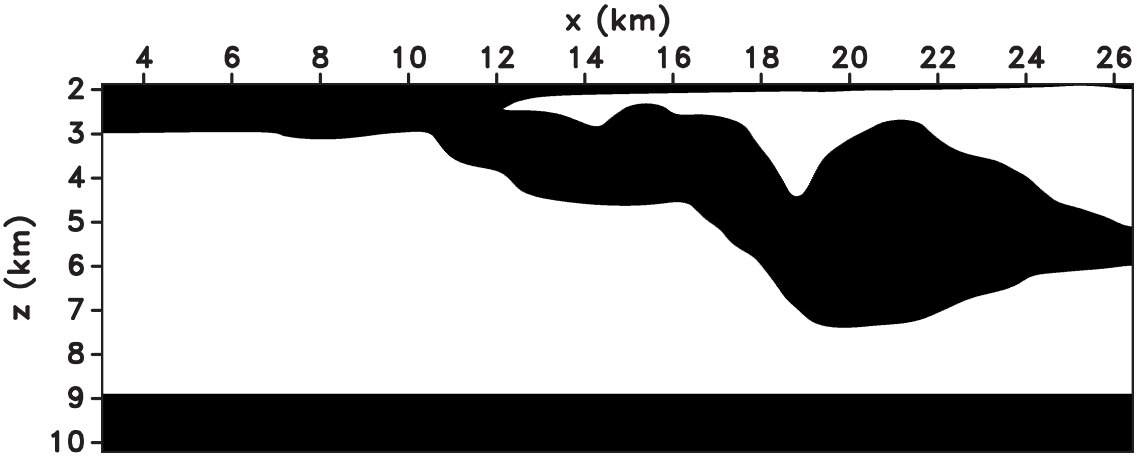


Figure 4.13. A mask of the salt and water bodies of the Sigsbee image.

planarity, the structure oriented semblance, and the envelope are allowed to constrain CIP locations.

4.9 Computing the Priority Map

As discussed at the beginning of this chapter, the priority map is computed by multiplying together the measurements discussed above. For 2D images, the linearity, the structure-oriented semblance, the envelope, and the mask are multiplied together. For 3D images, the planarity, the structure-oriented semblance, the envelope, and the mask are multiplied together. Figure 4.15 shows the PMAP constructed for the Sigsbee image, and Figure 4.16 shows the PMAP constructed for the SEAM image. Notice that the brightest values for both PMAPs typically correspond to the centers of coherent reflectors that have simple structure. Furthermore, values that are near faults, diffractors, or reflector truncations are either zero or small.

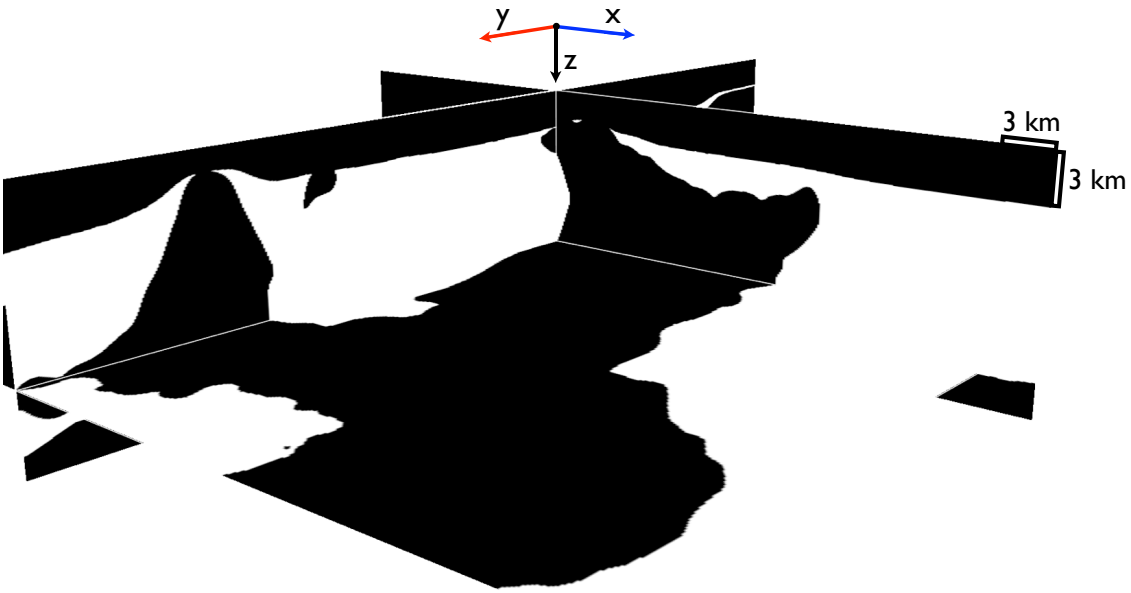


Figure 4.14. A mask of the salt and water bodies of the SEAM image.

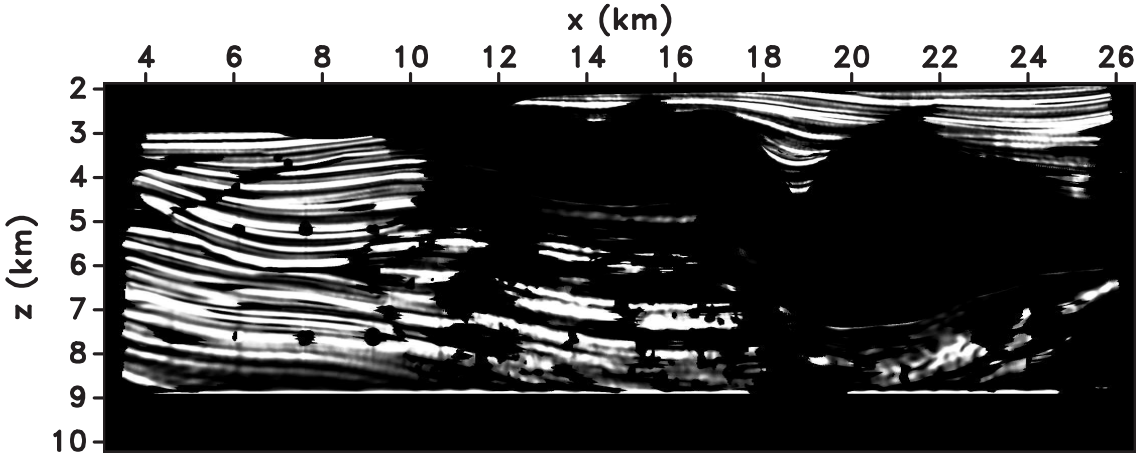


Figure 4.15. A priority map of the Sigsbee image. Notice that bright values correspond to locations that are near the center of coherent reflectors with simple structure.

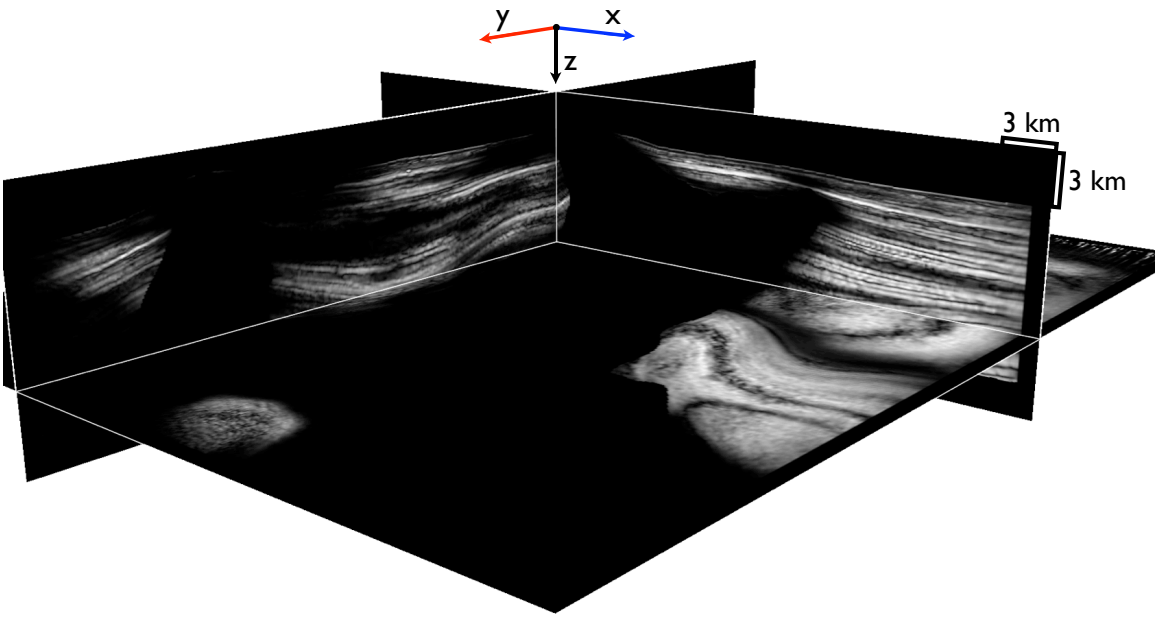


Figure 4.16. A priority map of the SEAM image. The bright values correspond to locations that are near the center of coherent reflectors with simple structure.

Chapter 5

Greedy Heuristic for Picking CIP Locations

To pick CIP locations based on the PMAP, or both the PMAP and local image structure, a greedy heuristic is used. However, before discussing this heuristic, the concept of a *greedy algorithm* and the difference between an algorithm and a heuristic are explained below.

5.1 Greedy Algorithms

Greedy algorithms are typically used to solve optimization problems such as the *fractional knapsack* problem. When solving these types of problems, a greedy algorithm "always makes the choice that looks best at the moment." (Cormen *et al.* , 2001) For example, the fractional knapsack problem is typically framed as a thief who is in the process of robbing some goods from a store. The thief has a knapsack that can only hold so much weight, and the goal of the thief is to take the most valuable load of goods as possible. The thief is able to take any fraction of a good that is in the store (e.g., pieces of bread, cuts of meat, handfuls of grain, etc.). In this case, the greedy strategy (the choice that looks best at the moment) is used by the thief to try to steal the most valuable load of goods. The greedy strategy that thief uses is to first take as much of the good with the greatest value-per-weight as possible. If the knapsack can still hold more goods, then the thief takes as much of the good with the next greatest value-per-weight as possible, and so forth until the knapsack cannot hold any more goods. It turns out that this strategy is an algorithm because it guarantees

that the thief leaves with the most valuable load of goods (Cormen *et al.* , 2001).

In contrast to the fractional knapsack problem, when the greedy strategy discussed above is applied to the *0–1 knapsack* problem, it doesn't guarantee the optimal solution (Cormen *et al.* , 2001). Therefore, the thief may not leave with the most valuable load of goods, and the greedy strategy in this case is a heuristic not an algorithm. The only difference between the two knapsack problems is that in the 0–1 knapsack problem, the thief can only take a whole good (e.g. a loaf of bread, a whole chicken, etc.).

5.2 Greedy Heuristic for Picking CIP Locations

As previously discussed, a greedy picking heuristic is used to pick CIP locations based on the PMAP or both the PMAP and local image structure. The objective of the greedy picker is to place CIPs at sparse locations throughout an image in manner that tries to optimize the sum of the PMAP values at the picked locations. The greedy strategy that is used by the picker is a sorting of the PMAP values into monotonically decreasing order. The greedy picker starts by picking the location that has the greatest PMAP value. Next, any location that is within a region (exclusion zone) that surrounds this pick is excluded from being picked in the future. This process is repeated until each location has either been picked, excluded, or has a PMAP value that is equal to zero.

Picked and excluded locations are tracked through the use of a data structure called a Point-of-Interest (POI). A POI is constructed for each location in the PMAP (i.e. for every pixel or voxel in the image). Each POI contains fields which hold its index position and its corresponding PMAP value. Furthermore, each POI has a boolean field called a *lock* which is set to *true* when the POI corresponds to a location that has been picked, excluded, or has a zero PMAP value.

For 3D images, the pseudocode for the greedy picking heuristic is shown in Heuristic 1. The heuristic for 2D images is the same except that the size of the third dimension n_3 is set to one. The array P_l is a linear array of POIs that have been constructed for every

location in the PMAP. The length of P_l is equal to $n_1 n_2 n_3$, where n_1, n_2, and n_3 are the size of the image and PMAP dimensions. The SORT procedure in Heuristic 1 is used to sort P_l in non-increasing order based on each POI's *value* field. The EXCLUDE-AT procedure in Heuristic 1 is used to excluded POIs that are within the exclusion zone of the most recent pick. This procedure takes as parameters the P_{3D} array, the *poi* that corresponds to the most recent pick, and the type of exclusion zone *ez* that is specified by the user. (Exclusion zones are discussed in the following section.) The P_{3D} array has the same dimensions as the image, and it contains the same POIs that are in P_l. Both P_{3D} and *poi* are need by the exclusion zone *ez* to determine which locations should be excluded based on the most recent pick.

Heuristic 1 Greedy CIP picker

```
1: SORT (P_l)
2: for j = 0 to < n_1 n_2 n_3 do
3:    if P_l[j].value = 0 then
4:       P_l[j].lock ← true
5:    end if
6: end for
7: i = 0
8: poi ← P_l[i]
9: while poi.lock ≠ true and i < n_1 n_2 n_3 do
10:   EXCLUDE-AT (P_3D,poi,ez)
11:   i ← i + 1
12:   poi ← P_l[i]
13: end while
```

5.3 Exclusion Zones

Exclusion zones are used to enforce sparseness between CIP locations, and these can be either isotropic or anisotropic regions. For isotropic exclusion zones (IEZs), these regions are circles in 2D and spheres in 3D. For anisotropic exclusion zones (AEZs), these regions are non-Euclidean ellipses in 2D and non-Euclidean ellipsoids in 3D. The shapes of the

AEZs semi-conform to the structure of local image features. The method used to define and construct AEZs is discussed below.

If IEZs are used in the greedy picking heuristic, then a picked CIP location cannot lie within the IEZ of another CIP location. However, if AEZs are used, this restriction is not always true, but if a CIP location does lie within the AEZ of another CIP, its placement will typically be near the boundary of the AEZ. The advantage of using AEZs is that CIP locations tend to be more sparse along structures directions then across them.

The radius of an IEZ is a user-defined parameter, and an example of a 2D IEZ is shown in Figure 5.2. This IEZ was computed for the Sigsbee image and it has a radius $r = 60$ pixels in z direction (x is undersampled compared to z so the radius is equivalent to the distance that corresponds to 60 pixels in z). Figure 5.3 and 5.4 show an example of a 3D IEZ computed for the SEAM image. This IEZ has a radius $r = 21$ voxels in the z direction (x and y are undersampled compared to z).

Determining the size and shape of an AEZ is more complicated than it is for an IEZ. To construct AEZs, a metric tensor field $\mathbf{D}(\mathbf{x})$ and an eikonal equation are used. The metric tensor field is derived from a structure tensor field $\mathbf{T}(\mathbf{x})$ that is computed for an image and a user-defined parameter which controls the non-Euclidean eccentricity of the AEZs. For each AEZ created at a picked location $\mathbf{x}_0$, a metric tensor field $\mathbf{D}(\mathbf{x})$ is used in an eikonal equation to solve for a measure of distance $t(\mathbf{x})$ between $\mathbf{x}_0$ and all points $\mathbf{x}$ that are within a distance t_{max} (in pixels or voxels) from $\mathbf{x}_0$. Therefore, the surface of each AEZ is defined by the points $\mathbf{x}$ such that $t(\mathbf{x}) = t_{max}$, where t_{max} is a user-defined parameter. The distances $t(\mathbf{x})$ are numerically solved using a fast iterative method that is similar to a method described by Jeong *et al.* (2007) and an eikonal equation described by Hale (2009a). This eikonal equation is defined as

$$\nabla t(\mathbf{x}) \cdot \mathbf{D}(\mathbf{x}) \nabla t(\mathbf{x}) = 1 \ , \tag{5.1}$$

where $t(\mathbf{x}_0) = 0$ is the boundary condition and the distance $t(\mathbf{x})$ is a non-Euclidean distance.

For a 2D AEZ, the metric tensor $\mathbf{D}$ is defined as

$$\mathbf{D} = e_u \mathbf{u}\mathbf{u}^T + \mathbf{v}\mathbf{v}^T \ , \tag{5.2}$$

and for a 3D AEZ, the metric tensor $\mathbf{D}$ is defined as

$$\mathbf{D} = e_u \mathbf{u}\mathbf{u}^T + \mathbf{v}\mathbf{v}^T + \mathbf{w}\mathbf{w}^T \ , \tag{5.3}$$

where e_u is the non-Euclidean eccentricity parameter and $0 < e_u \leq 1$. When $e_u = 1$, an AEZ is isotropic and equivalent to an IEZ.

Each time the greedy heuristic picks a location $\mathbf{x}_0$, the origin of (5.1) is set to $\mathbf{x}_0$, and all locations $\mathbf{x}$ that have not already been picked or excluded and that are a distance $t(\mathbf{x}) < t_{max}$ are then excluded from being picked in the future.

Figure 5.2 shows an example of a 2D AEZ computed for the Sigsbee image. The white region is the AEZ for a single location. Notice that the AEZ conforms to the structure of the local reflectors(s). For this AEZ, $t_{max} = 60$ pixels and $e_u = 0.1$. If t_{max} were decreased, then the AEZ shown would be smaller in both the structure and cross-structure directions. If e_u were decreased, then the AEZ would be thinner only in the cross-structure direction, and if increased, the AEZ would be thicker in this direction.

An example of a 3D AEZ is shown in Figure 5.5 and 5.6. For this AEZ, $t_{max} = 20$ voxels and $e_u = 0.05$. As in the 2D case, the 3D AEZ conforms to the structure of the local reflectors. Furthermore, if t_{max} were decreased then the AEZ would be smaller in all directions. If e_u were decreased, then the AEZ would be thinner only in the cross-structure direction, and if increased, the AEZ would be thicker in this direction.

5.4 Runtime Complexity

The runtime complexity of the greedy heuristic depends on the type of exclusion zones used. When IEZs are used, the runtime complexity of the heuristic is $O(nlog(n))$. However,

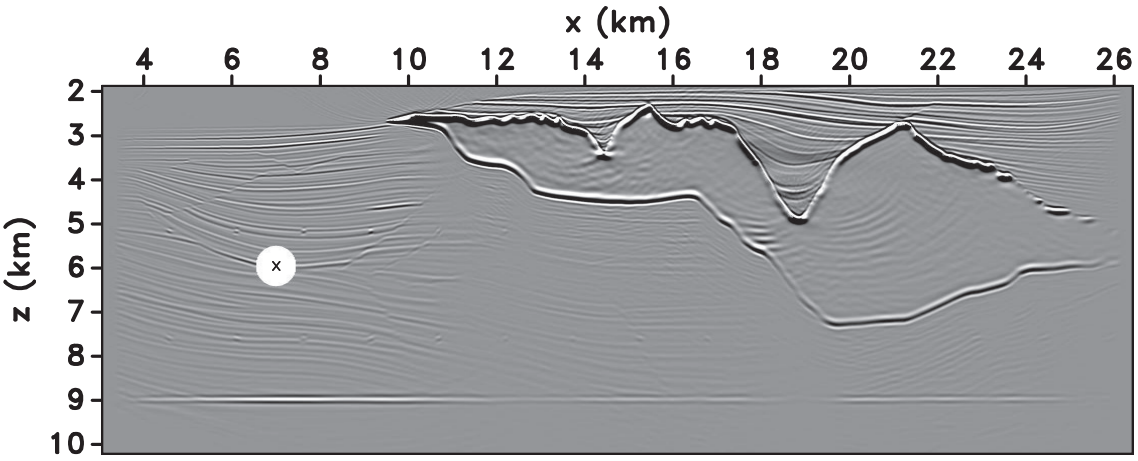


Figure 5.1. The Sigsbee image with an isotropic exclusion zone.

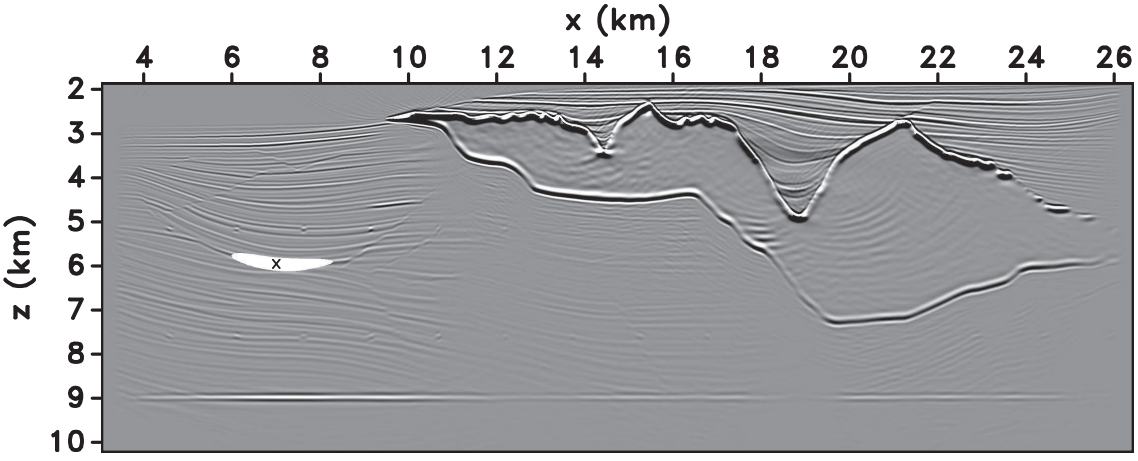


Figure 5.2. The Sigsbee image with an anisotropic exclusion zone.

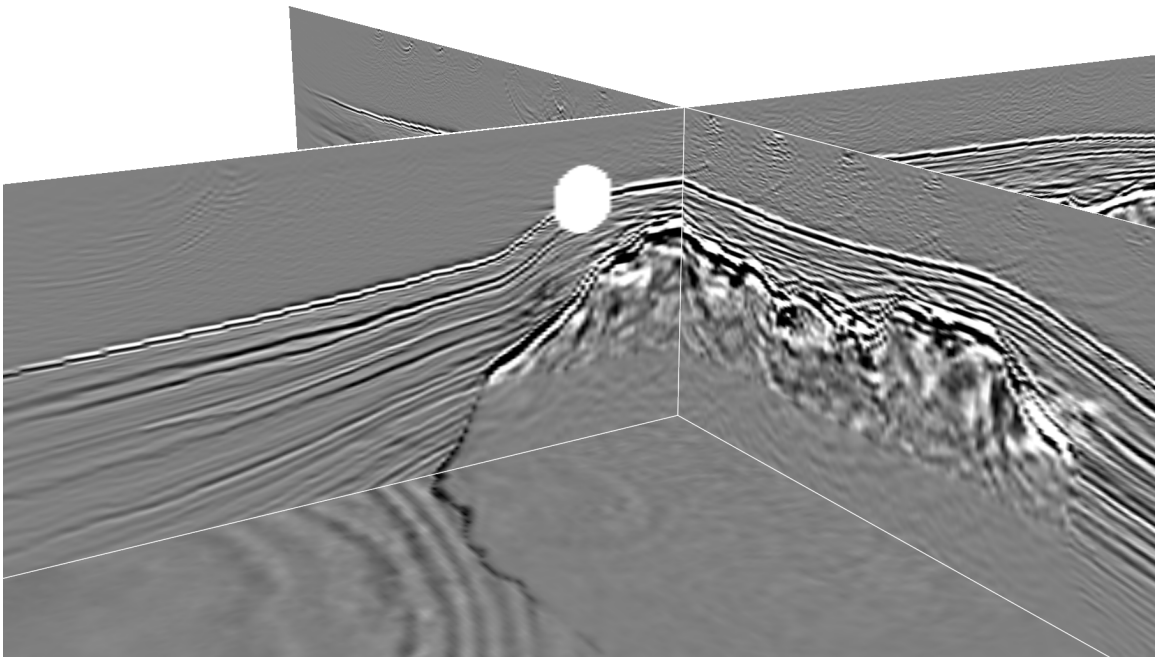


Figure 5.3. The SEAM image with an isotropic exclusion zone.

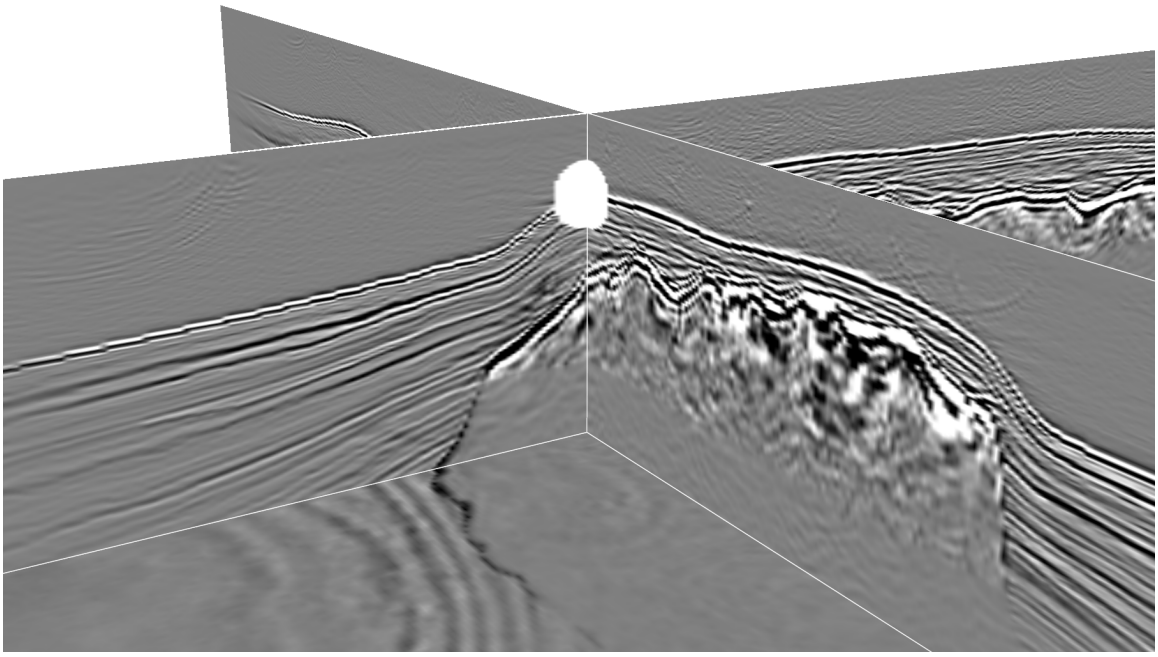


Figure 5.4. Another view of the the isotropic exclusion zone shown in Figure 5.4.

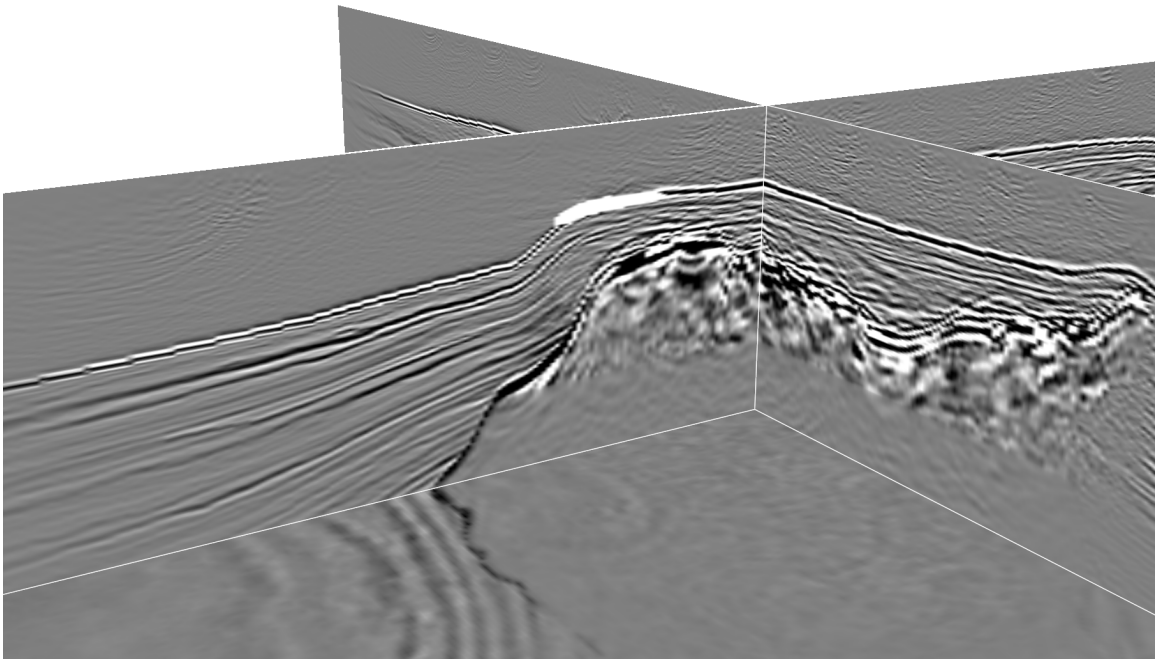


Figure 5.5. The SEAM image with an anisotropic exclusion zone.

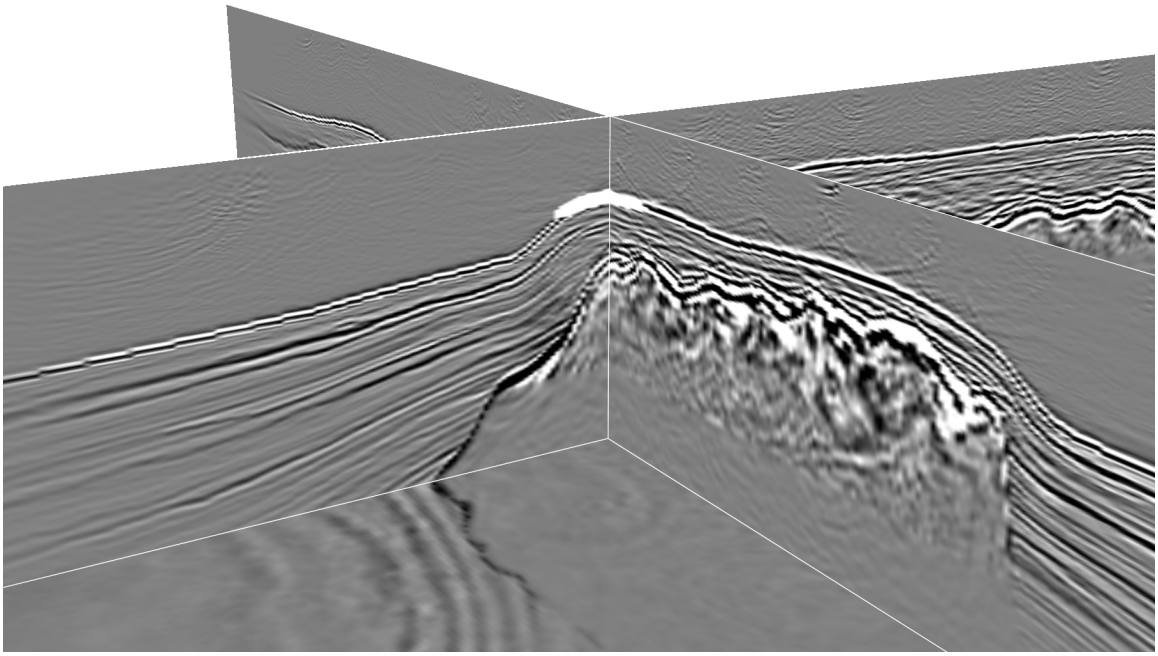


Figure 5.6. Another view of the anisotropic exclusion zone shown in Figure 5.6.

when AEZs are used, then the runtime complexity depends on both the size of the exclusion zones and the number of iterations needed to solve the eikonal equation (5.1).

When IEZs are used, the while-loop in the greedy heuristic has a runtime complexity of $O(cm+n)$, where n is the number of pixels or voxels in an image, c is the maximum number of pixels or voxels that compose an exclusion zone, and m is the number of CIP locations picked. However, $c \propto \frac{n}{m}$ because an exclusion zone is the size of at least a single pixel or voxel and at most the size of the image (for a pick that is located in a corner of the image). If $c = 1$, then $m = n$ because every location gets picked, and if $c = n$, then $m = 1$ because only one location can be picked. Furthermore, when $1 < c < m$, the maximum average overlap between an exclusion zone and its nearest neighbors is limited due to geometry. Therefore, the while-loop has a runtime that is $O(cm + n) = O(n + n) = O(n)$. The initialization of the P_l and P_{3D} arrays used in the greedy picking heuristic has a runtime complexity of $O(n)$, and the runtime complexity of the SORT procedure has a runtime of $O(nlog(n))$. The remaining operations in the heuristic have a runtime of $O(1)$. Because the SORT procedure has the highest order runtime complexity, the runtime of the greedy picking heuristic is $O(nlog(n))$ when IEZ's are used.

As previously stated, when AEZs are used in the greedy heuristic, the runtime complexity depends on the size of the exclusion zones and the number of iterations needed to solve the eikonal equation. However, in practice, the CPU runtime of the heuristic will be significantly less than the time it would take to manually pick CIP locations. For example, the picks shown in Figure 6.5 and 6.6 were computed in less than 10 minutes using two 2.66 GHz Quad-Core Intel Zeon processors (note the greedy heuristic is a serial process). Furthermore, the priority map (shown in 4.16) that was used to generate the picks shown in Figure 6.5 and 6.6 was computed in less than 15 minutes using the same quad-core processors.

Chapter 6

Results

Using the PMAPs that were created for the Sigsbee and SEAM images, the picks for CIP locations were computed using both isotropic exclusion zones and anisotropic exclusion zones.

6.1 Sigsbee 2A Data

Figure 6.1 shows the picks that were selected for the Sigsbee image using IEZs that have a radius $r = 60$ pixels in z. Notice that the majority of the picks lie along the center of coherent reflectors that have simple structure. Furthermore, the picks sample the image somewhat uniformly. The picks shown in Figure 6.2 were selected using AEZs with a maximum distance $t_{max} = 40$ pixels and an eccentricity $e_u = 0.2$. Notice that the majority of the picks also lie along the center of coherent reflectors of simple structure. However, in contrast to the picks computed using IEZs, the picks computed using AEZs do not sample the image uniformly, although the image is still well sampled. Instead, these picks tend to be more sparse along the reflectors and more dense across them.

6.2 SEAM Data

The picks shown in Figure 6.3 and 6.4 were selected for the SEAM image using IEZs with a radius $r = 21$ voxels in z. Both of these figures show the same results (the top 5000 picks from a total of 28,018) but from different viewpoints. The picks were plotted with the

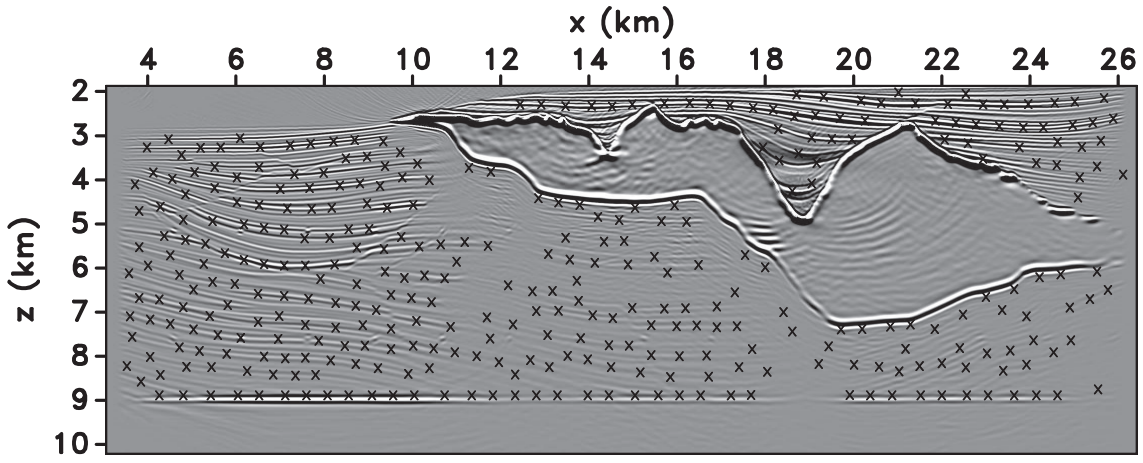


Figure 6.1. The Sigsbee image with picks using IEZs.

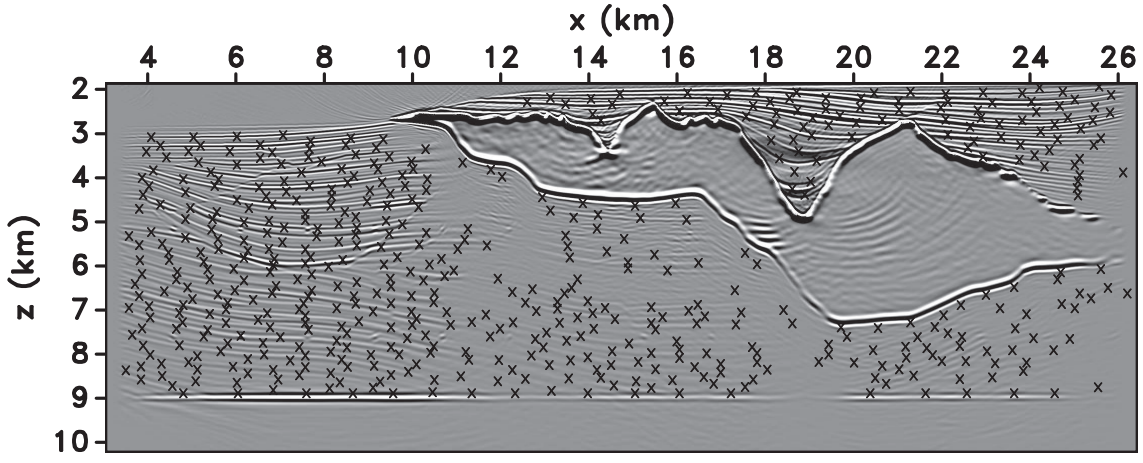


Figure 6.2. The Sigsbee image with picks using AEZs.

PMAP instead of the image because this makes it is easier to see the picks in 3D. However, visualizing the quality of these picks using just Figure 6.3 and 6.4 is difficult. This difficulty is because these figures are 2D projections of a 3D image. However, using software, a user can interact with the results by rotating, translating, and zooming in on the image, as well as sliding the 2D panels back and forth and up and down. Using these techniques, a user can verify that the majority of the picks lie along the center of coherent reflectors that have simple structure, and that the picks sample the image somewhat uniformly. The results shown in Figure 6.3 and 6.4 were verified using software that was developed from the Mines Java-Tool-Kit.

Figure 6.5 and 6.6 show the top 5000 picks for the SEAM image from a total of 18,218 picks. These picks were selected using AEZs with a maximum distance $t_{max} = 15$ voxels and an eccentricity $e_u = 0.34$. Both of these figures show the same results but from different viewpoints. Again, the picks were plotted with the PMAP to assist with visualizing the results. Using the techniques discussed above, a user can verify that the majority of the picks lie along the center of coherent reflectors that have simple structure. Furthermore, a user can verify that although the picks do not sample the image uniformly, the image is still well sampled. As in the 2D case, these picks tend to be more sparse along the reflectors and more dense across them.

6.3 Verifying the Quality of 3D Picks

As discussed above, visually verifying the quality of the picks selected for 3D images is more difficult than it is for 2D images. For 2D images, verifying is easily done by plotting the picks on top of the image or the PMAP. However, for 3D images, a user needs to interact with the results using software. To assist with the quality verification of 3D picks, tensor ellipsoids can be plotted at each picked location. The method used to compute these ellipsoids is described by Engelsma & Hale (2009). These ellipsoids conform to the local image structure both in shape and in orientation. If a pick is located within a locally

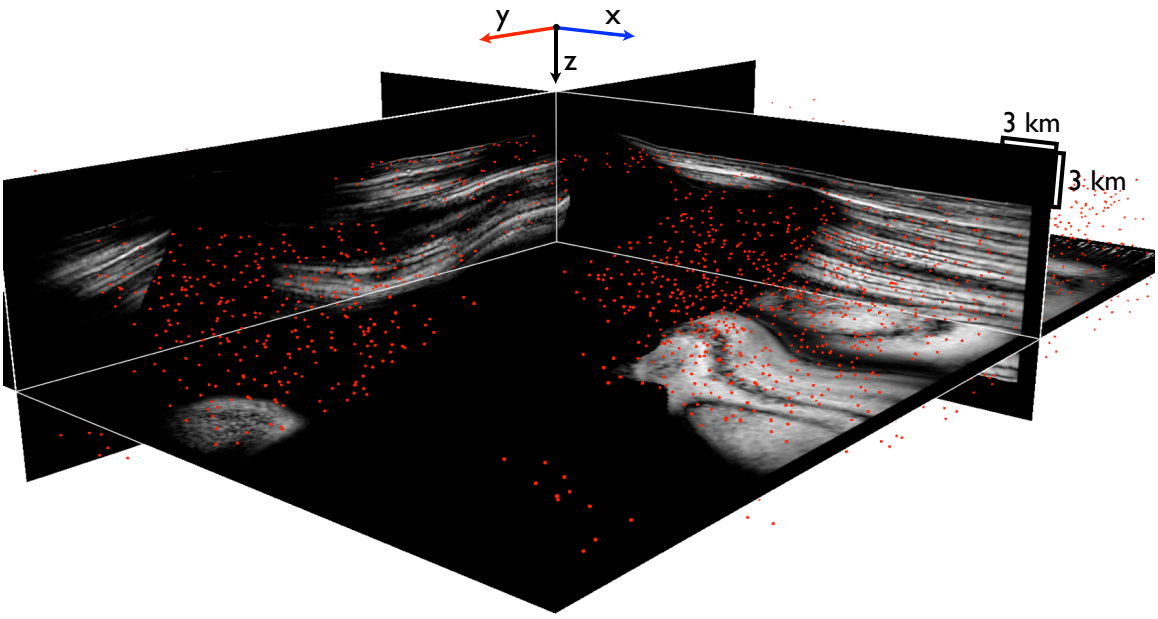


Figure 6.3. The SEAM image with CIP locations that were picked using IEZs. The top 5000 picks are shown.

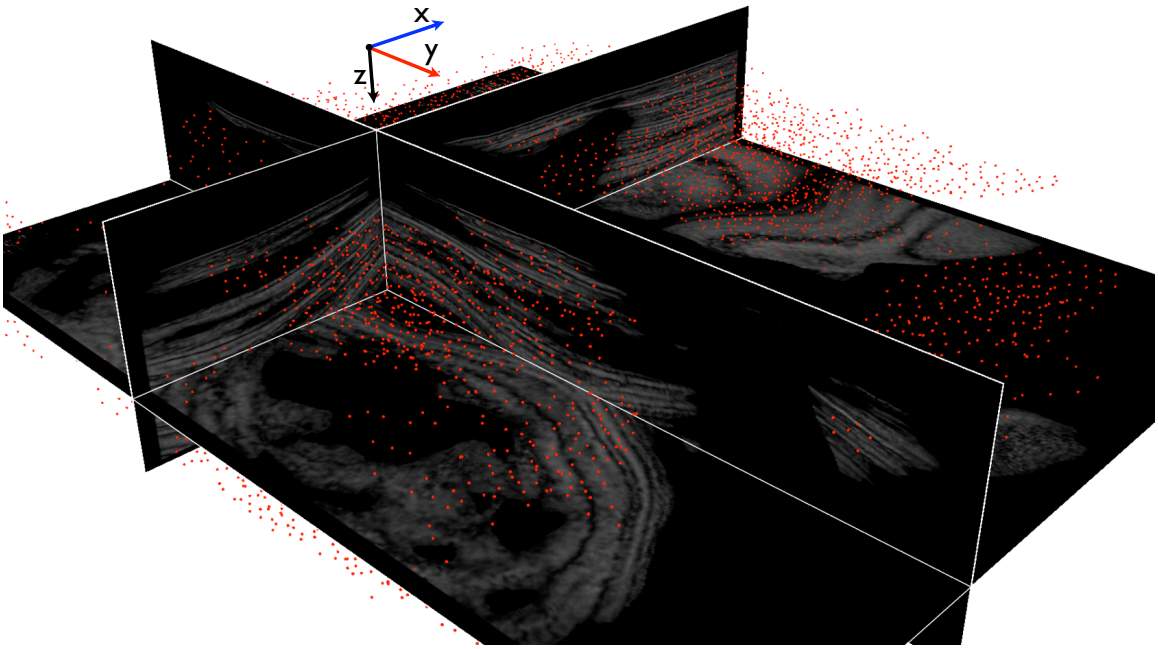


Figure 6.4. Another view of the SEAM image with CIP locations that were picked using IEZs. The top 5000 picks are shown.

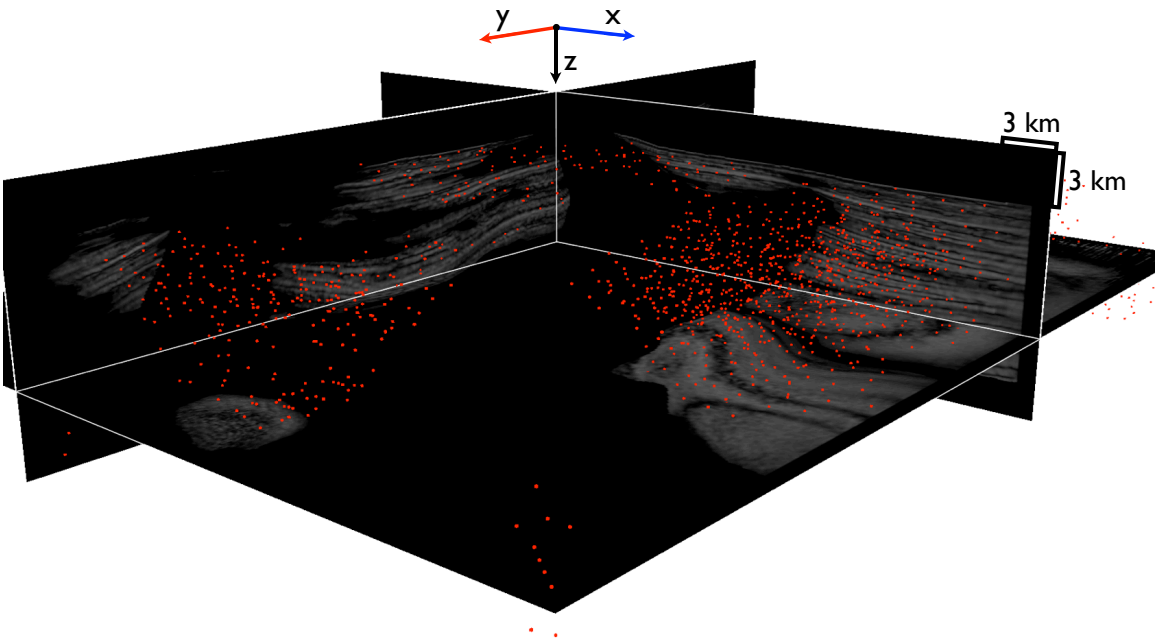


Figure 6.5. The SEAM image with CIP locations that were picked using AEZs. The top 5000 picks are shown.

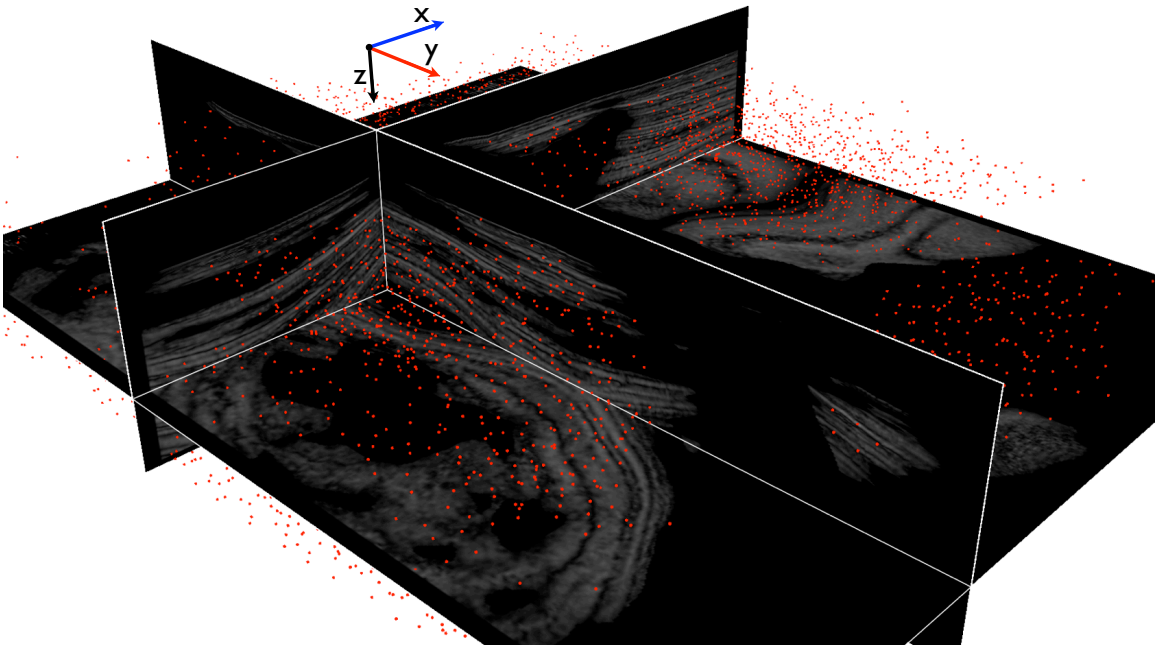
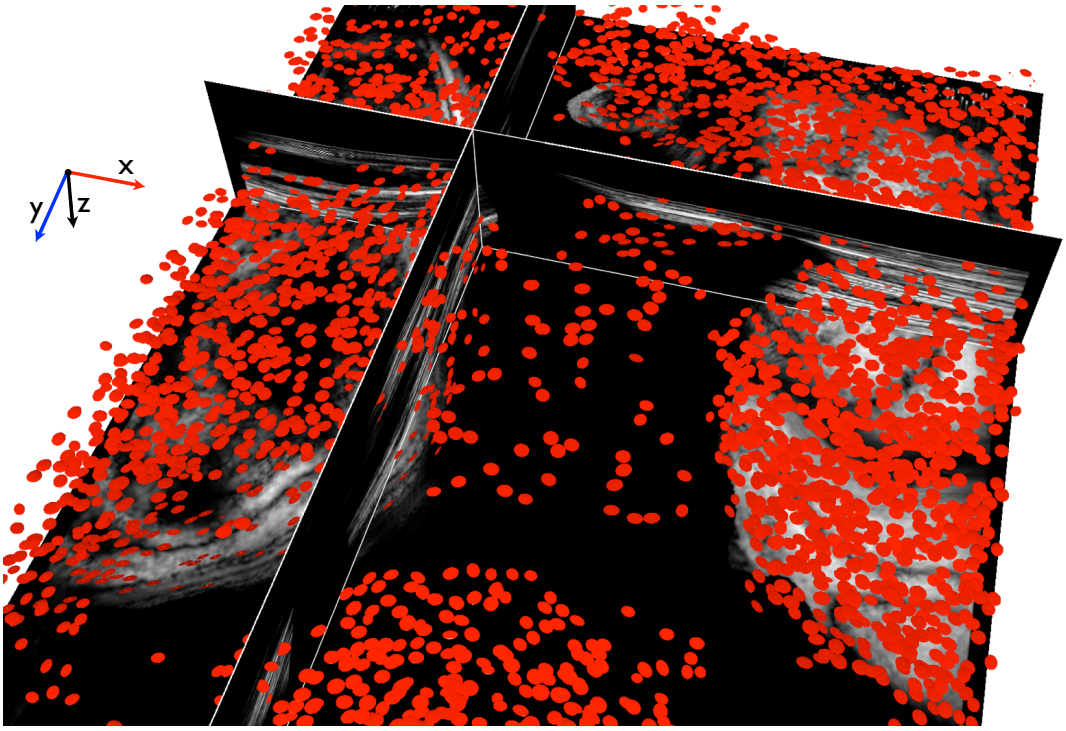


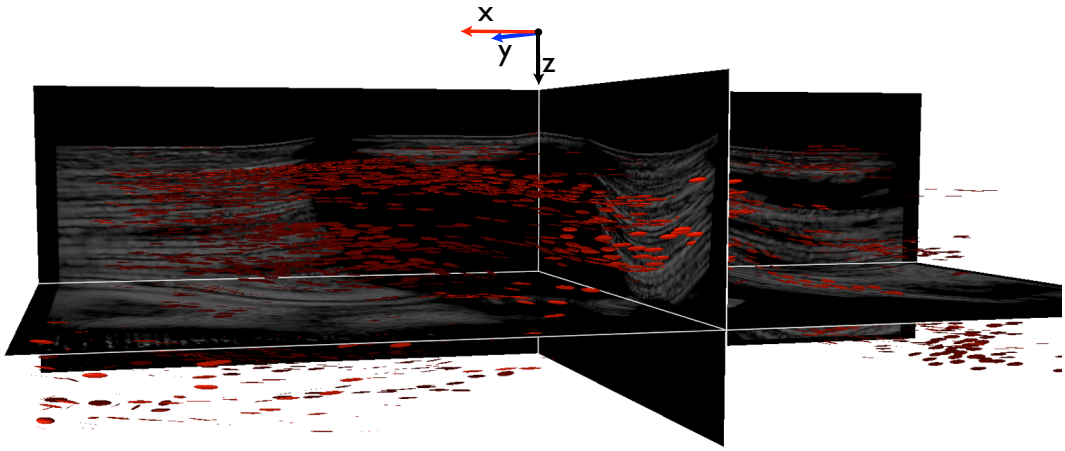
Figure 6.6. Another view of the SEAM image with CIP locations that were picked using AEZs. The top 5000 picks are shown.

voluminous feature, then the ellipsoid at this location be mostly spherical. If a pick is located within a locally planar reflector (e.g. a geologic unit), then the corresponding ellipsoid will be disk shaped. For thick/thin reflectors, this disk will be somewhat thick/thin. If a pick lies within a curvilinear reflector (e.g. a river or stream channel), then the corresponding ellipsoid will be elongated and narrow. Finally, if a pick lies within a locally complex feature, then the ellipsoid will have a shape that is a mix between between spherical, disklike, and elongated.

Ideally, the ellipsoids that correspond to picked locations should be disk shaped because the PMAP should constrain the picks to simple, coherent, planar reflectors. Figure 6.7(a) and 6.7(b) show the tensor ellipsoids plotted for the picks shown in Figure 6.3 and 6.4. A bird's-eye view of these ellipsoids is shown in Figure 6.7(a), and a side view is shown in Figure 6.7(b). Notice that the overwhelming majority of the ellipsoids are disk-shaped, and that their orientations (dip and strike) correspond to the orientations of the local image features.



(a)



(b)

Figure 6.7. (a) A bird's-eye view of the PMAP of the SEAM image with structure conforming ellipsoids, and (b) a side view of the PMAP with the structure conforming ellipsoids.

Chapter 7

Conclusions

An automated method for picking CIPs from seismic images is described and demonstrated in this thesis. This method has two stages. In the first stage, a priority map is constructed which constrains the placement of CIPs to image locations that are near the center of coherent reflectors that have simple structure. The PMAP is constructed by multiplying together measurable image properties and an optional user-defined mask. These image properties include linearity, planarity, structure-oriented semblance, and the amplitude envelope. The linearity and planarity measurements provide a measure of structure complexity, and they are derived from structure-tensors. Structure-tensors are also used with a method that computes the structure-oriented semblance, which provides a measure of structural complexity and coherency. The amplitude envelope provides a measure of a location's nearness to the center of a reflector, and the mask is used to manually exclude the placement of CIPs within specific regions of an image such as within salt or water bodies.

In the second stage, a fast greedy picking heuristic is used to automatically pick CIP locations. The PMAP or both the PMAP and local image structures are used to guide the placement of CIP locations. Exclusion zones are used to control the sparseness between picked locations, and they can be either isotropic exclusion zones which are disks in 2D and spheres in 3D or anisotropic exclusion zones which are non-Euclidean ellipses in 2D and non-Euclidean ellipsoids in 3D. If IEZs are used in the greedy heuristic, then picked CIP locations will sample an image somewhat uniformly. If AEZs are used, then CIP locations

will tend to be more sparse along reflectors, but the image will still be well sampled. The greedy picker also maintains an ordering of the picks based on their corresponding PMAP values, and because of this ordering, it is easy to remove lower quality picks. The quality of the picks in 2D can be visually verified by plotting the picks on top of the image or the PMAP. For 3D images, the quality of the picks can be visually verified by interacting with software and by displaying the tensor ellipsoids that correspond to picked CIP locations.

To be most effective for MVA, CIPs need to be constructed at sparse locations throughout an image and near the center of coherent reflectors that have simple structure. However, manually picking sparse CIP locations from complex seismic images can be a difficult and time-consuming process. The method described in this thesis is both fast and effective at automatically picking sparse CIP locations locations that are guided by measurable image properties.

References

Al-Yahya, Kamal. 1989. Velocity analysis by iterative profile migration. *Geophysics*, **54**(6), 718–729.

Berkhout, A. J., & Pao, Y.-H. 1982. Seismic Migration–Imaging of Acoustic Energy by Wave Field Extrapolation. *Journal of Applied Mechanics*, **49**(3), 682–683.

Claerbout, Jon F. 1985. *Imaging the earth's interior*. Cambridge, MA: Blackwell Scientific Publications, Inc.

Clapp, Robert G. 2001. Ray-based tomography with limited picking. *SEP Report 110*, 103–113.

Cormen, Thomas H., *et al.* . 2001. *Introduction to Algorithms*. 2nd edn. Cambridge, MA: McGraw-Hill Higher Education.

Engelsma, Chris, & Hale, Dave. 2009. Visualization of 3D tensor fields derived from seismic images. *CWP Report 655*, 187–192.

Fehmers, Gijs C., & Hocker, Christian F. W. 2003. Fast structural interpretation with structure-oriented filtering. *Geophysics*, **68**(4), 1286.

Gray, Samuel H., Etgen, John, Dellinger, Joe, & Whitmore, Dan. 2001. Seismic migration problems and solutions. *Geophysics*, **66**(5), 1622.

Hale, Dave. 2009a. Image-guided blended neighbor interpolation of scattered data. *SEG Technical Program Expanded Abstracts*, **28**(1), 1127–1131.

Hale, Dave. 2009b. Structure-oriented smoothing and semblance. *CWP Report 635*, 261–270.

Jeong, Won-Ki, Fletcher, P. Thomas, Tao, Ran, & Whitaker, Ross. 2007. Interactive Visualization of Volumetric White Matter Connectivity in DT-MRI Using a Parallel-Hardware Hamilton-Jacobi Solver. *IEEE Transactions on Visualization and Computer Graphics*, **13**, 1480–1487.

Sava, Paul, & Vasconcelos, Ivan. 2011. Extended imaging conditions for wave-equation migration. *Geophysical Prospecting*, **57**(1), 35–55.

Shen, Peng, Symes, William W., & Stolk, Christiaan C. 2003. Differential semblance velocity analysis by wave-equation migration. *SEG Technical Program Expanded Abstracts*, **22**(1), 2132–2135.

Stolk, Christiaan C., de Hoop, Maarten V., & Symes, William W. 2005. Kinematics of prestack shot-geophone migration. *SEG Technical Program Expanded Abstracts*, **24**(1), 1866–1869.

Symes, William W. 2008. Migration velocity analysis and waveform inversion. *Geophysical Prospecting*, **56**(6), 765–790.

Vliet, Lucas J. Van, & Verbeek, Piet W. 1995. Estimators for Orientation and Anisotropy in Digitized Images. *Pages 442–450 of: ASCI95, Proc. First Annual Conference of the Advanced School for Computing and Imaging (Heijen, NL, May 16-18), ASCI.*

Yang, Tongning, & Sava, Paul. 2010. Wave-equation migration velocity analysis with extended commom-image-point gathers. *CWP Report 649*, 101–120.