

Chapter 2

Time-lapse seismic cross-equalization

Time-lapse (4D) seismic images provide important information about production- or injection-related changes in reservoir properties. However, in practice, imperfect repetition of survey geometries, non-repeatable ambient noise, changes in natural environmental conditions, and other discrepancies generate artifacts that contaminate production-related amplitude differences between seismic images. Before reservoir property changes can be extracted from time-lapse seismic images, these artifacts must be attenuated, while production-related differences must be preserved.

In conventional time-lapse seismic imaging, the process of attenuating unwanted differences between seismic data sets is called cross-equalization or cross-matching (Rickett and Lumley, 2001; Hall et al., 2005). Together with better-quality data derived from improved acquisition technology, in many cases, seismic cross-equalization makes it possible to extract reliable information about reservoir property changes from time-lapse seismic images. However, as discussed in later chapters, seismic cross-equalization methods are inapplicable or inadequate in several practical scenarios. Before introducing new methods that overcome some limitations of seismic

cross-equalization in later chapters, it is important to first consider conventional processing methods and how they can be improved.

Depending on the monitoring goals and the data quality, seismic cross-equalization methods can be applied at different stages of the processing sequence—before and/or after imaging. In general, it is desired that throughout the processing sequence, identical parameters and steps are applied to individual data sets. Prior to imaging, some key preprocessing steps include data regularization, noise and multiple attenuation, amplitude and phase corrections, and trace selection (4D-binning). These preprocessing steps are well developed and are applied in routine time-lapse processing (Helgerud et al., 2011b; Sharma et al., 2011). As discussed in later chapters, careful data preprocessing is also important for the inversion methods developed in this dissertation.

In this chapter, I consider two common post-imaging cross-equalization steps: warping and matched filtering. Warping involves estimating the apparent displacements (misalignments) between images, which are caused by production-induced velocity changes and deformation. Such apparent displacements can be used to estimate velocity and geomechanical changes between surveys (Hatchell and Bourne, 2005b; Røste et al., 2006; Hawkins et al., 2007). In addition, prior to computing the amplitude difference between baseline and monitor images, the computed apparent displacements are used to align the images, thereby correcting for compaction and velocity errors.

Using a method adapted from Hale (2009), I show how multidimensional warping can be performed efficiently with a sequential one-dimensional (1D) approach. This approach prevents errors in time-lapse images commonly caused by approximating multidimensional warping by a 1D vertical warping; it also avoids the computational expense of full multidimensional cross-correlation, search and interpolation. To obtain accurate information about production-related changes from time-lapse seismic images, all displacement components must be considered. Importantly, this warping method will be applied in later chapters as a preprocessing step for the inversion methods developed in this chapter.

Matched filtering removes residual amplitude and phase differences between the aligned images. First, matched filters are designed in a region outside the reservoir, where no change is expected. These matched filters are then applied to the full data, including the reservoir. I show how optimal matched filters can be estimated with evolutionary programming and that matched filters estimated this way provide more reliable time-lapse images than conventional methods. Combining this matched filtering method with the sequential warping method described above provides an efficient and robust cross-equalization tool.

In the following sections, I first discuss the sequential warping and optimized matched filtering methods. I then apply these methods to two time-lapse seismic data sets from the Norne field. I show that these cross-equalization methods provide a robust way to obtain good quality time-lapse images.

MULTIDIMENSIONAL WARPING BY SEQUENTIAL 1D CROSS-CORRELATION

Reservoir depletion and fluid injection cause stress-induced compaction and stretching in and around a producing reservoir, which in turn cause changes in seismic travel times and path lengths between surveys. However, relative to the background velocity and reflector depth, the resulting velocity and geomechanical changes between surveys are typically small. Furthermore, because these velocity and geomechanical changes are coupled, it is often difficult to estimate accurate migration velocities for individual data sets without introducing additional errors that may mask the production-induced changes of interest. Therefore, in practice, baseline and monitor data sets are usually imaged with a single (baseline) velocity model, leading to amplitude and kinematic differences between the baseline and monitor images. It is assumed that relative to amplitude differences caused by reservoir fluid and rock property changes, small errors introduced by imaging the monitor data with the baseline velocity are negligible.

To correctly estimate production- or injection-related time-lapse amplitudes, kinematic differences, which show up as *apparent* displacements between the images, can

be corrected by multidimensional warping (Rickett and Lumley, 2001; Hall et al., 2005). However, because it is expensive to estimate all components of the apparent displacement field, and because of the generally horizontal layering of sedimentary units, it is common practice to consider only vertical components. Instead of this *vertical-only* approximation, it is possible to efficiently estimate all displacement components using a sequential 1D warping approach (Hale, 2009). As shown by Hale (2009), this sequential 1D warping approach produces components of apparent displacement vectors that are equivalent to results from full multidimensional warping. Other approaches to solving the multidimensional warping problem as a combination of several sequential 1D warping problems have been published by other authors (Aarre, 2008). In order to obtain accurate estimates of reservoir property changes, all displacement components must be taken into account.

Here, to implement a multidimensional warping operator, I follow the approach of Hale (2009). This method, which only requires sequential 1D searches for cross-correlation peaks along different axes, is summarized in Algorithm 1. Given two seismic traces b and m , the normalized correlation coefficient C_{bm} , computed within a sliding window of half-width h around sample point k , is given by

$$C_{bm}(l) = \frac{\sum_{i=k-h}^{k+h} b(i) \cdot m(i+l)}{\sqrt{\sum_{i=k-h}^{k+h} b(i)^2} \cdot \sqrt{\sum_{i=k-h}^{k+h} m(i+l)^2}}, \quad (2.1)$$

where l denotes the correlation lag between samples in the traces. For the time-lapse problem, b and m represent baseline and monitor data, respectively. The cross-correlation at each sample point is localized by a Gaussian taper W given by

$$W(j) = e^{-\frac{j^2}{2r^2}}, \quad (2.2)$$

where q is the sample distance relative to point k , and r is the Gaussian width. In this dissertation, after testing different values of r , I find that 1/4 the half-width of the correlation window gives a satisfactory result.

Algorithm 1 Apparent displacements by cyclic 1D search

```

b  $\leftarrow$  preprocess( $b_o$ )
m  $\leftarrow$  preprocess( $m_o$ )
while  $s^{ia=1:3} \geq s_{stop}^{ia=1:3}; j = 1 : niter$  do
  for axis = 1:3 do
    if axis = 1 then
       $s^1 = s[xcorr_{peak}(m_j|b)]$  !compute displacement components along axis 1
       $z^1 \leftarrow s^1 < z^1 > + s^1$  !interpolate/accumulate components along axis 1
       $z^2 \leftarrow s^1 < z^2 >$  !interpolate along components along axis 2
       $z^3 \leftarrow s^1 < z^3 >$  !interpolate along components along axis 3
       $m_j^1 \leftarrow s^1 < m_j >$  !interpolate monitor along components along axis 1
    else if axis = 2 then
       $s^2 = s[xcorr_{peak}(m_j^1|b)]$ 
       $z^1 \leftarrow s^2 < z^1 >$ 
       $z^2 \leftarrow s^2 < z^2 > + s^2$ 
       $z^3 \leftarrow s^2 < z^3 >$ 
       $m_j^2 \leftarrow s^2 < m_j^1 >$ 
    else if axis = 3 then
       $s^3 = s[xcorr_{peak}(m_j^2|b)]$ 
       $z^1 \leftarrow s^3 < z^1 >$ 
       $z^2 \leftarrow s^3 < z^2 >$ 
       $z^3 \leftarrow s^3 < z^3 > + s^3$ 
       $m_j^3 \leftarrow s^3 < m_j^2 >$ 
       $m_{j+1} \leftarrow m_j^3$ 
    end if
  end for
end while
 $m_o \leftarrow z^{1:3} < m_o >$ 

```

j – iteration;

ia – axis;

b_o – original baseline data

m_o – original monitor data

b – preprocessed baseline data

m – preprocessed monitor data

$xcorr_{peak}$: correlation peak

s^{ia} – apparent displacement along axis ia

z^{ia} – cumulative displacements along axis ia

$w < m >$: interpolation of m with displacements w

Although this implementation closely follows the method proposed in Hale (2009), I introduce several additional constraints that further improve the method. The most important changes are summarized below:

- To improve the resolution of the correlation peaks along the horizontal axes, I preprocess the input images by removing near-zero wavenumber components along these axes. Along any axis, if the data within a cross-correlation window are dominated by near-zero wavenumbers, it becomes impossible to reliably estimate the apparent displacements between them.
- To first capture the lower-frequency apparent displacement components, this iterative procedure starts with large correlation gates and lags. Then, to capture higher-frequency components, the correlation gates and lags are systematically decreased as a function of iteration.
- To improve the accuracy of the estimated displacements, prior to searching for the cross-correlation peaks, I smooth the correlation coefficients by averaging the values from neighboring traces. In addition, I define an acceptance criterion based on the relative cross-correlation peaks and apparent displacements such that outliers are discarded. For example, at points in the image volume where the peak cross-correlation values are relatively small and the estimated apparent displacements relatively large, the displacement components are rejected and then replaced by values interpolated from neighboring points.

As shown in Algorithm 1, for each iteration, the results (accumulated displacement components and warped images) along one axis (e.g., vertical) serve as inputs for computation along the next axis (e.g., in-line). For each axis, displacement components are accumulated after interpolating the image and displacement components from the previous step. This accumulation procedure is explained in detail by Hale (2009). In this dissertation, I perform all interpolations with a cubic spline algorithm (Press et al., 1996).

Figure 2.1 shows two synthetic seismic images and the horizontal and vertical displacement components between them. These displacement components represent vertical downward and horizontal outward apparent displacements of image points in Figure 2.1(b) relative to those in Figure 2.1(a). The only difference between these two images is misalignments of image points described by these displacements (Figures 2.1(c) and 2.1(d)). To align these images, I apply the sequential 1D warping

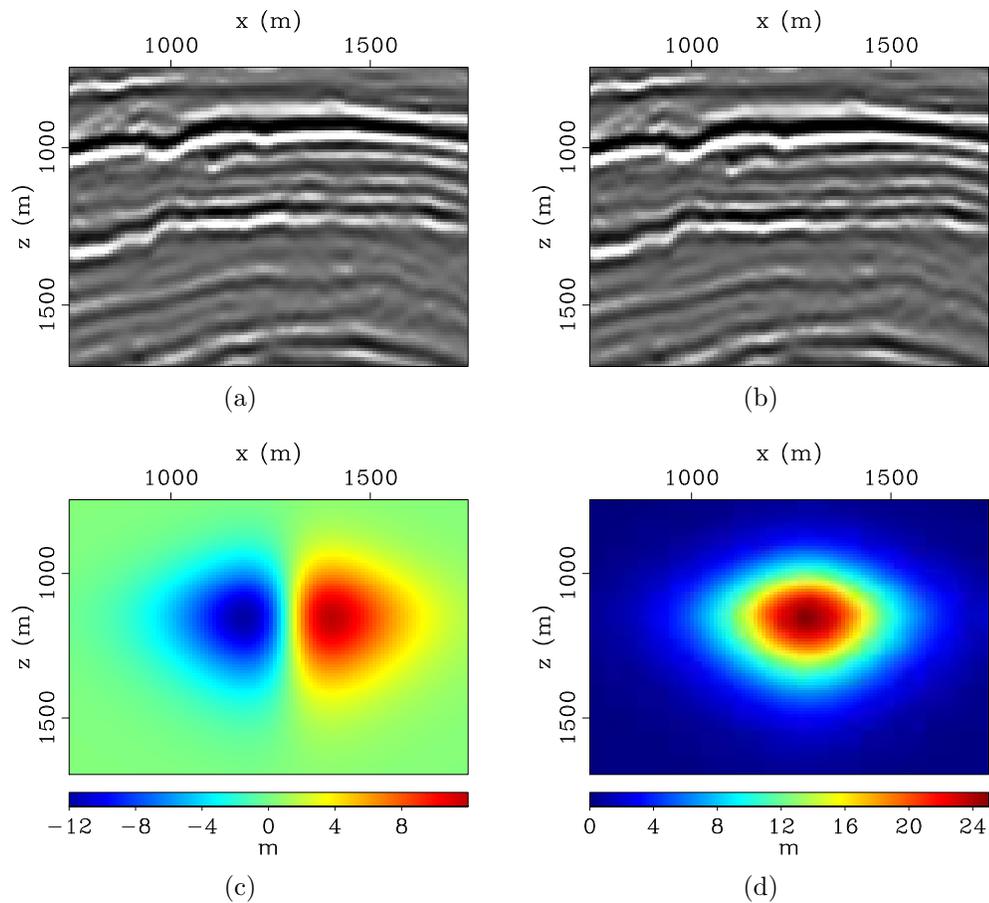


Figure 2.1: Synthetic seismic images (a) and (b), and components of the apparent displacements between them along the horizontal (c) and vertical (d) axes. The image differences before, and after different kinds of warping are shown in Figure 2.2. The evolution of apparent displacement components derived from sequential 1D warping are presented in Figures 2.3 and 2.4. [CR].

chap2/. s-warp-1,s-warp-2,s-warp-x,s-warp-z

method described above to them. For comparison, I also align these images using one iteration of vertical warping (neglecting all horizontal components), and ten iterations of vertical warping. The image differences before and after warping are presented in Figure 2.2.

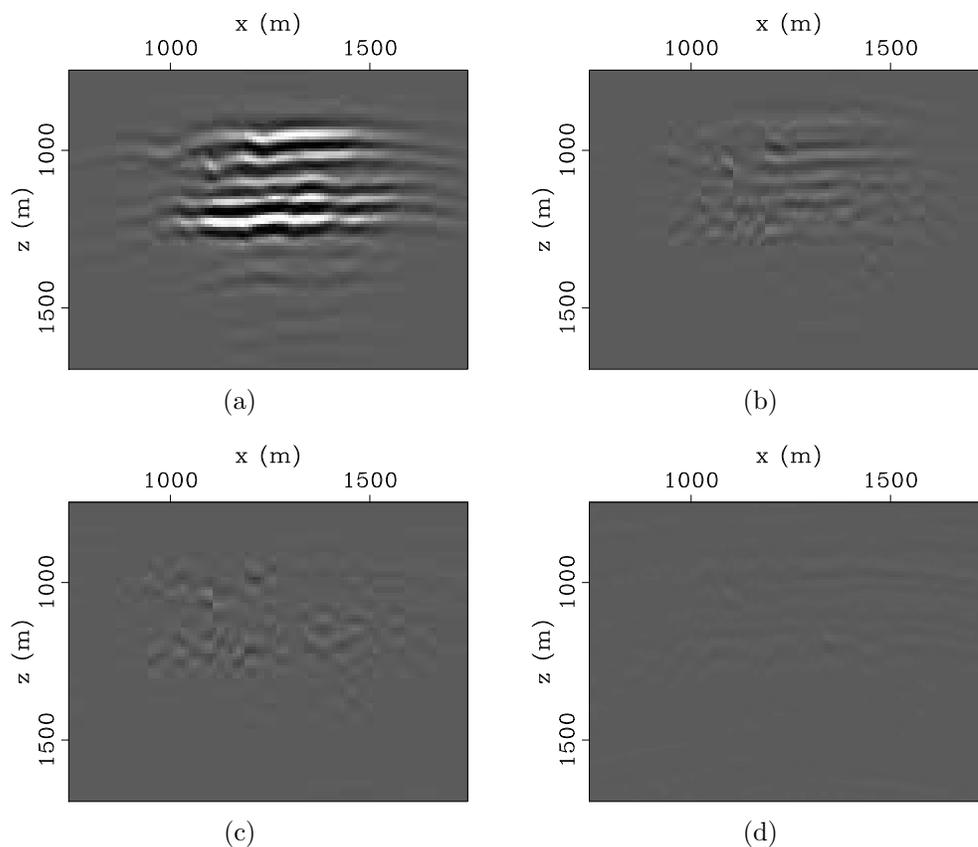


Figure 2.2: Differences between the seismic images in Figures 2.1(a) and 2.1(b) before warping (a), after one iteration of vertical warping (b), after 10 iterations of vertical warping (c), and after 10 iterations of sequential one-dimensional warping. All images are clipped to similar amplitude levels. Although iterative vertical warping provides improved match compared to one iteration of *conventional* vertical warping (b), there are still significant differences between the two images (c). Sequential 1D warping provides the best results with negligible amplitude differences between the warped images (d). [CR]. chap2/. s-warp-d-b4,s-warp-d-afO,s-warp-d-afz,s-warp-d-af

In Figure 2.2, because one iteration of vertical warping is a poor approximation to the multidimensional displacements between the images, it provides poor alignment of these images—hence the large residual differences between them (Figure 2.2(b)). In addition, although iterative vertical warping provides an improved alignment of the images, there are still large residual differences between the aligned images (Figure 2.2(c)). However, using the sequential 1D warping method, residual amplitude differences between the aligned images are negligible (Figure 2.2(d)). Negligible amplitude differences in Figure 2.2(d) suggests that this sequential 1D warping method can sufficiently resolve multidimensional apparent displacements between images. Figures 2.3 and 2.4 show how the two components of the apparent displacement vectors evolve with iterations. Note that as a function of iterations, estimates of the displacement components converge toward the actual displacements (Figures 2.1(c) and 2.1(d)).

As previously stated, the computed apparent displacements serve two general purposes. First, prior to computing the amplitude difference between the baseline and monitor images, these displacements are used to align the images, thereby ensuring that all events are collated. Furthermore, it is possible to obtain important information about changes within and around the reservoir from these apparent displacements. For example, the fractional velocity change $\frac{\delta v}{v}$ and vertical strain ϵ_{zz} can be computed as follows (Hatchell and Bourne, 2005a; Hale, 2009):

$$\frac{\delta v}{v} = -\frac{R}{1+R} \frac{d(\delta t)}{dt}, \quad (2.3)$$

$$\epsilon_{zz} = -\frac{1}{R} \frac{\delta v}{v}, \quad (2.4)$$

where v is the baseline velocity, δv is the change in velocity between the baseline and monitor surveys, $d(\delta t)$ is the corresponding vertical displacement (time-shift), and R is the dilation factor. For the examples considered in this chapter, I assume a constant dilation factor, $R = 5$, a value within the range measured for North Sea reservoirs (Hatchell and Bourne, 2005a).

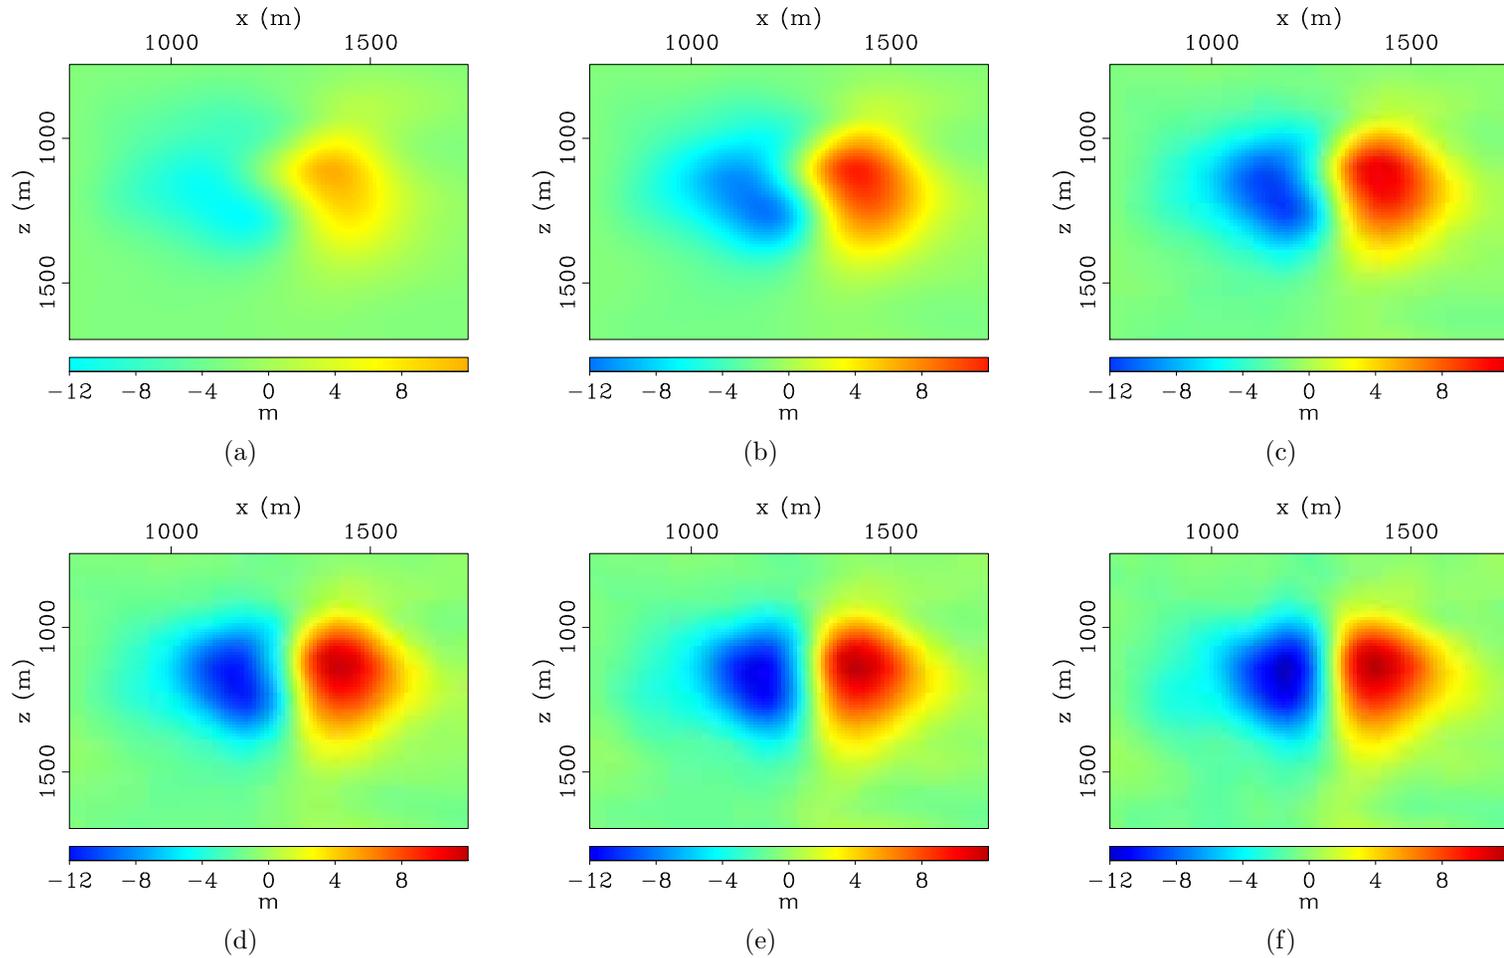


Figure 2.3: Evolution of the horizontal displacement components between Figures 2.1(a) and 2.1(b) obtained via sequential 1D warping. These panels show the estimated displacements after one, two, four, six, eight and ten iterations, respectively. As a function of iterations, these displacements converge towards the actual values in Figure 2.1(c). [CR]. `chap2/. s-warp-x1,s-warp-x2,s-warp-x3,s-warp-x4,s-warp-x5,s-warp-x6`

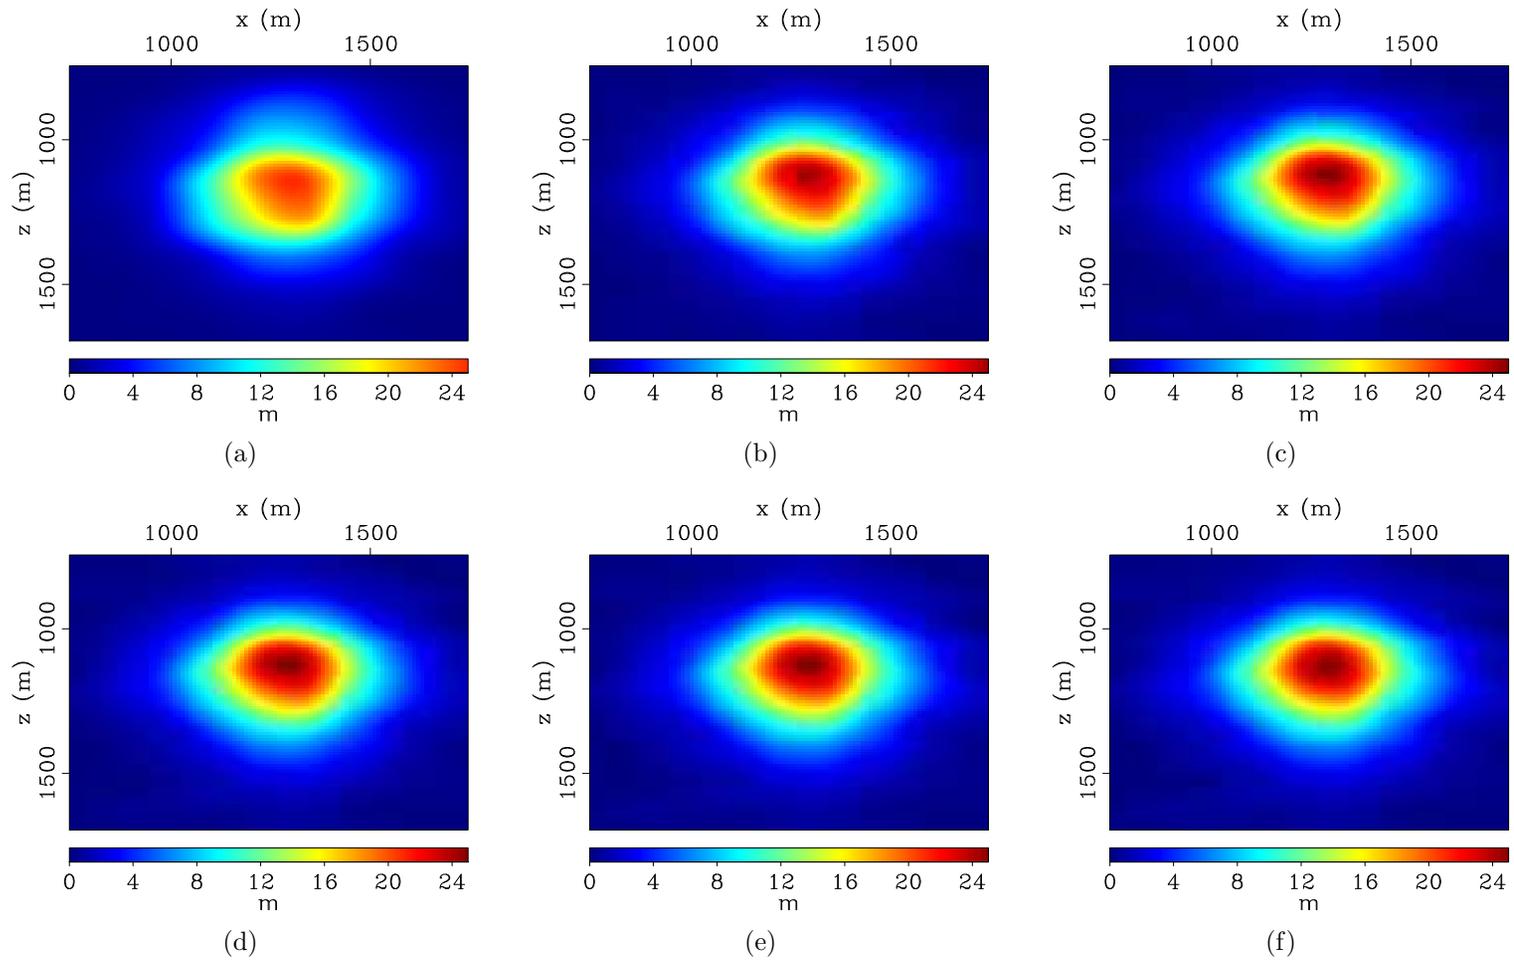


Figure 2.4: Evolution of the vertical displacement components between Figures 2.1(a) and 2.1(b) obtained via sequential 1D warping. These panels show the estimated displacements after one, two, four, six, eight and ten iterations, respectively. As a function of iterations, these displacements converge towards the actual values in Figure 2.1(d). [CR]. chap2/. s-warp-z1,s-warp-z2,s-warp-z3,s-warp-z4,s-warp-z5,s-warp-z6

MATCHED FILTERING BY EVOLUTIONARY PROGRAMMING

To ensure that only production-related amplitude changes within the reservoir are interpreted, it is usually desired that amplitude and phase differences in non-reservoir regions be minimal. In time-lapse cross-equalization, this requirement is usually achieved by matched filtering. In general, it is assumed that non-production-related amplitude and phase differences that contaminate non-reservoir regions in the image volume are similar to the ones that contaminate the reservoir region. Therefore, by applying a matched filter that minimizes amplitude and phase differences in a non-reservoir region to a reservoir region, non-production-related artifacts within the reservoir can be attenuated. The filtering strategy can be global (a single filter for all trace locations), local (different filters for different trace locations), or a combination of the two. In this section, I consider only local one-dimensional matched filters.

Given baseline and monitor data \mathbf{b} and \mathbf{m} , the filter \mathbf{f} that matches the two sets of data in a least-squares sense is one that minimizes the quadratic cost function r given by

$$r(\mathbf{f}) = \|\mathbf{M}\mathbf{f} - \mathbf{b}\|_2, \quad (2.5)$$

where \mathbf{M} is a convolution matrix constructed from coefficients of the monitor data.

The matched filter obtained by minimizing the cost-function in equation 2.5 satisfies the expression

$$\mathbf{f} = (\mathbf{M}^T\mathbf{M})^{-1}\mathbf{M}^T\mathbf{b}, \quad (2.6)$$

where the superscript T denotes the matrix transpose.

In the frequency domain, equation 2.6 becomes

$$\mathbf{F}(w) = \frac{\overline{\mathbf{M}(w)}\mathbf{B}(w)}{\overline{\mathbf{M}(w)}\mathbf{M}(w)}, \quad (2.7)$$

where $\mathbf{F}(w)$ is the matched filter, $\mathbf{B}(w)$ is the baseline data, and $\mathbf{M}(w)$ is the monitor data at frequency w . The overline denotes complex conjugation.

To avoid division by zero, a damping factor ϵ is included in the denominator of equation 2.7 so that it becomes

$$\mathbf{F}(w) = \frac{\overline{\mathbf{M}}(\omega)\mathbf{B}(\omega)}{\overline{\mathbf{M}}(\omega)\mathbf{M}(\omega) + \epsilon^2}. \quad (2.8)$$

As mentioned earlier, at each trace location, the matched filter is derived within a window outside the reservoir, where no change is expected, and then it is applied to the full trace data. However, the matched filter derived at each trace location is a function of the estimation parameters. In this implementation, these parameters include the filter length, the top and bottom boundaries of the estimation window, and the damping factor. Conventionally, a single set of parameters derived by a manual trial-and-error approach is used at all trace locations. However, because the set of parameters that produces satisfactory filters in some parts of the image volume may produce very poor filters in other parts, obtaining an optimal set of parameters can be a tedious or impossible challenge.

In this dissertation, I apply an evolutionary programming (EP) method to estimate the optimal filter parameters for all trace locations. Evolutionary programming belongs to a class of global optimization methods called evolutionary algorithms (Bäck, 1996; Baeck et al., 1997; Yao et al., 1999; Eiben and Smith, 2003). These algorithms solve optimization problems using Darwinian evolutionary principles of natural selection (see Algorithm 2). In this application, I use the EP method because of its flexibility and robustness. Further discussions application of this method to selecting matched filtering parameters are presented in Ayeni and Nasser (2009) and Ayeni (2011).

The most important implementation considerations for the EP algorithm include the population size and initialization, the mutation operator, and the selection/rejection criterion. For the examples presented in the next section, I find that a population size of 20 produces satisfactory results. The initial population is derived as random parameters within $\pm 10\%$ bounds of the best set of predetermined parameters. During each iteration, the *fittest* individuals in the population are perturbed to generate the

offspring (i.e., new sets of parameters) for the next iteration, whereas the worst half are rejected. The fittest individuals (half the population size) are defined as those with the lowest predefined error/fitness function.

Selecting an appropriate fitness function is critical, because this determines which of the solutions are kept and which are rejected at each iteration. In this dissertation, the fitness function is defined as the normalized root-mean-squared (NRMS) difference within a non-reservoir *validation* window which is different from that used to estimate the filter. The NRMS difference N_{bm} within the validation window is defined as

$$N_{bm} = 2 \times \frac{\sqrt{\sum_{i=1}^n b(i) - m(i)}}{\sqrt{\sum_{i=1}^n b(i) + \sum_{i=1}^n m(i)}}, \quad (2.9)$$

where n is the number of samples within the validation window.

By selection matched filter parameters such that NRMS difference within a different non-reservoir window is minimized, I ensure that no new artifacts are introduced into the matched data. As I shown in the field data example, one disadvantage of matched filtering as a cross-equalization tool is that if the parameters are poorly choosing, production-related amplitude differences can be contaminated by spurious matched filtering artifacts. Selecting filtering parameters that satisfy predefined criteria—such as minimum energy within a validation window—reduce such filtering spurious artifacts.

To ensure continuity from trace to trace, a smoothness constraint is applied and the estimation procedure repeated with an initial population based on the smoothed parameters. Furthermore, an additional smoothness constraint ensures that filter coefficients vary slowly from trace to trace. In the case where several partial stacks are to be matched, to preserve time-lapse Amplitude Variation with Angle (AVA) information, a further constraint can be introduced to ensure that the matched filters vary slowly between partial stacks.

Algorithm 2 Optimal matched filtering by evolutionary programming

```

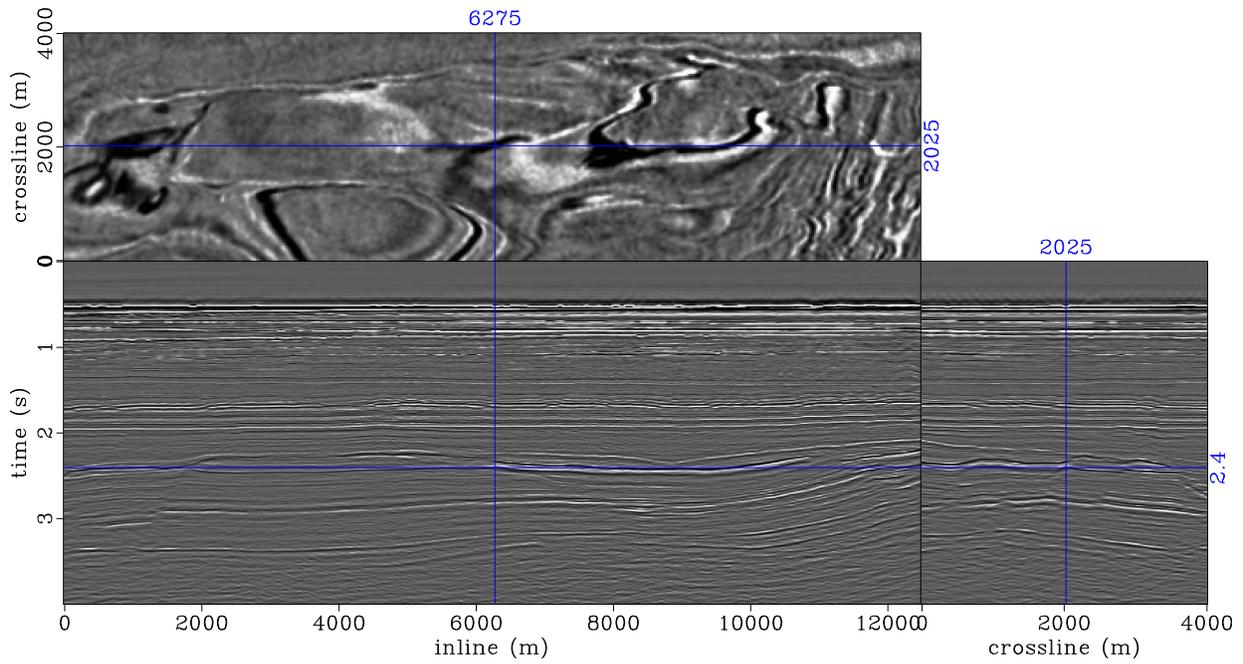
for  $irepeat = 1:nrepeat$  do
  for  $itrace = 1:ntrace$  do
    (Re)Initialize: estimation window, filter length, etc.
    Evaluate: fitness (nrms)
    while  $nrms_{iter} \geq nrms_{stop}; iter = 1 : niter$  do
      Select: fittest individuals
      Mutate: generate new offspring
      Evaluate: fitness (nrms)
    end while
  end for
  Apply constraints
end for

```

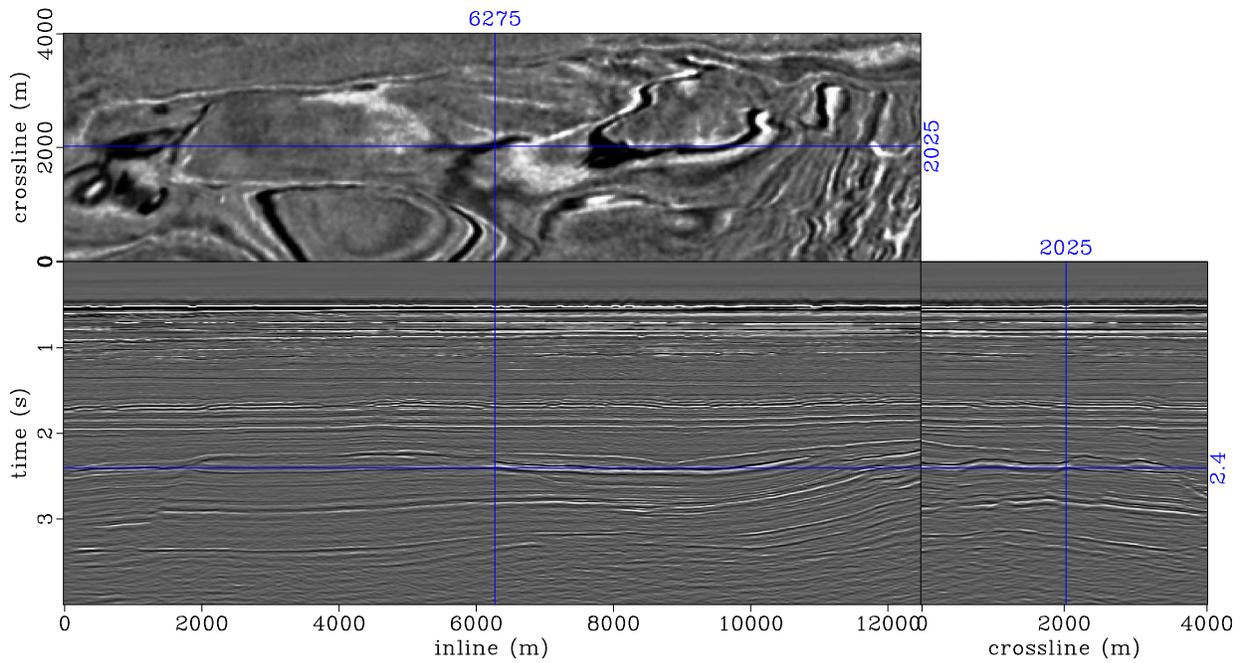
FIELD DATA APPLICATION

In this section, I apply the warping and matched filtering methods described in the previous sections to four data sets from the Norne field, which is located in the Norwegian North Sea. The main field is a 9 km x 3 km horst block composed of high-porosity, high-permeability, high net-to-gross lower and middle Jurassic sandstones (Osdal et al., 2006). The field was discovered in 1991 and production started in 1997.

The baseline data were acquired in 2001, and the three monitor data sets were acquired in 2003, 2004 and 2006. These data sets have been preprocessed using conventional methods and imaged with Kirchhoff prestack time migration. Figure 2.5 shows full-stack images from the baseline and 2006 monitor data. Osdal et al. (2006) and Aarre (2008) provide detailed descriptions of the field geology, its production history, seismic data acquisition and processing, and some interpretation of the observed time-lapse seismic amplitudes.



(a)



(b)

Figure 2.5: Migrated images of the Norne field obtained from the 2001 baseline (a), and 2006 Monitor (b) data sets. The data processing steps are discussed by Osdal et al. (2006). [CR]. `chap2/. mig-full-2001,mig-full-2006`

First, using the sequential warping method described in a previous section, I estimate the apparent displacements between the baseline and the monitor images. For each sample point, at the first iteration, I use correlation windows of 0.4 s two-way travel-time, and 500 m half widths along the inline and crossline axes. These are systematically reduced to half the starting values at the third and final iteration. As discussed above, an acceptance criterion is defined such that points with relatively large displacements and low peak cross-correlation values are discarded. Specifically, in this example, image points with correlation coefficients lower than the 5th percentile and apparent displacements greater than the 95th percentile are discarded and interpolated from neighboring traces.

Figure 2.6 shows the vertical components of the apparent displacements (i.e., time-shifts) between the baseline and the 2006 monitor images after different iterations. Note that as mentioned earlier, the resolution of the displacements increases with the number of iterations. In Figure 2.6 (and in similar displays throughout this dissertation), the three panels are slices through a 3D volume at the positions indicated by the orthogonal lines. These time-shifts clearly show parts of the Norne horst block that have been produced between 2001 and 2006.

The lateral—inline and crossline—components of the apparent displacements between the baseline image and the 2006 monitor image are presented in Figures 2.7(a) and 2.7(b). Neglecting these lateral displacement components will create errors in the measured time-shifts (Figure 2.7(c)). In the examples presented in this chapter, I define errors in time-lapse attributes (time-shifts, velocity change, amplitudes, etc.) as the difference between values derived from considering only vertical displacement components and those derived from considering vertical and lateral displacement components. The time-shifts and lateral displacements extracted along the top of the reservoir are presented in Figure 2.8. In this figure, note that apparent displacements provide qualitative information about changes in properties around the reservoir. Quantitative interpretation of these displacement components is outside the scope of the current study.

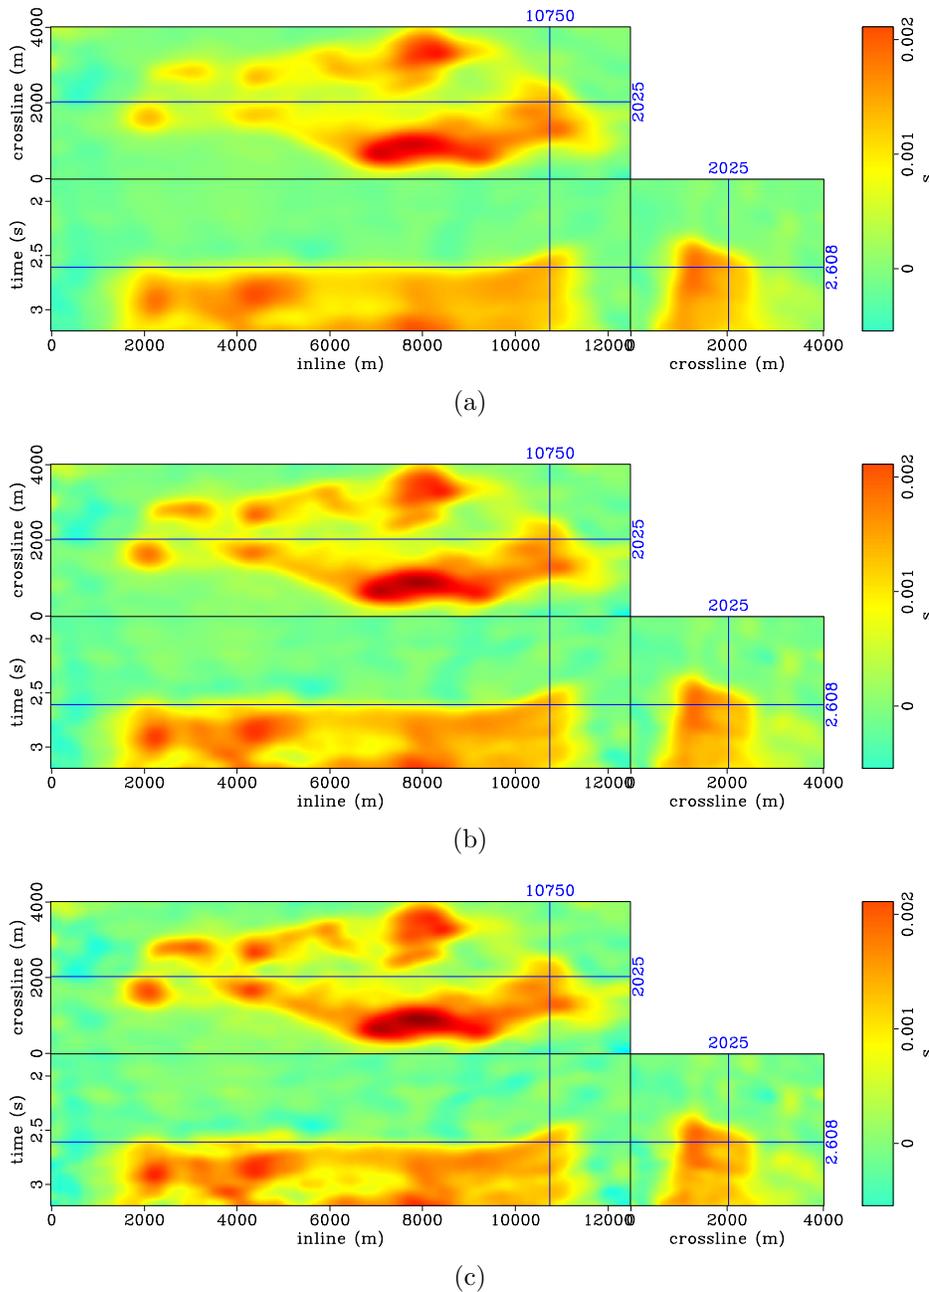


Figure 2.6: Vertical displacement components (time-shifts) between the 2006 monitor image and the baseline after one (a), two (b) and three (c) iterations. Note that the first iteration captures main features of the displacement field, whereas later iterations capture the higher-frequency details, thereby improving the resolution. The top of the producing segments between 2.5 s and 2.6 s can be mapped easily using these displacement components. [CR]. `chap2/. ts-20061,ts-20062,ts-20063`

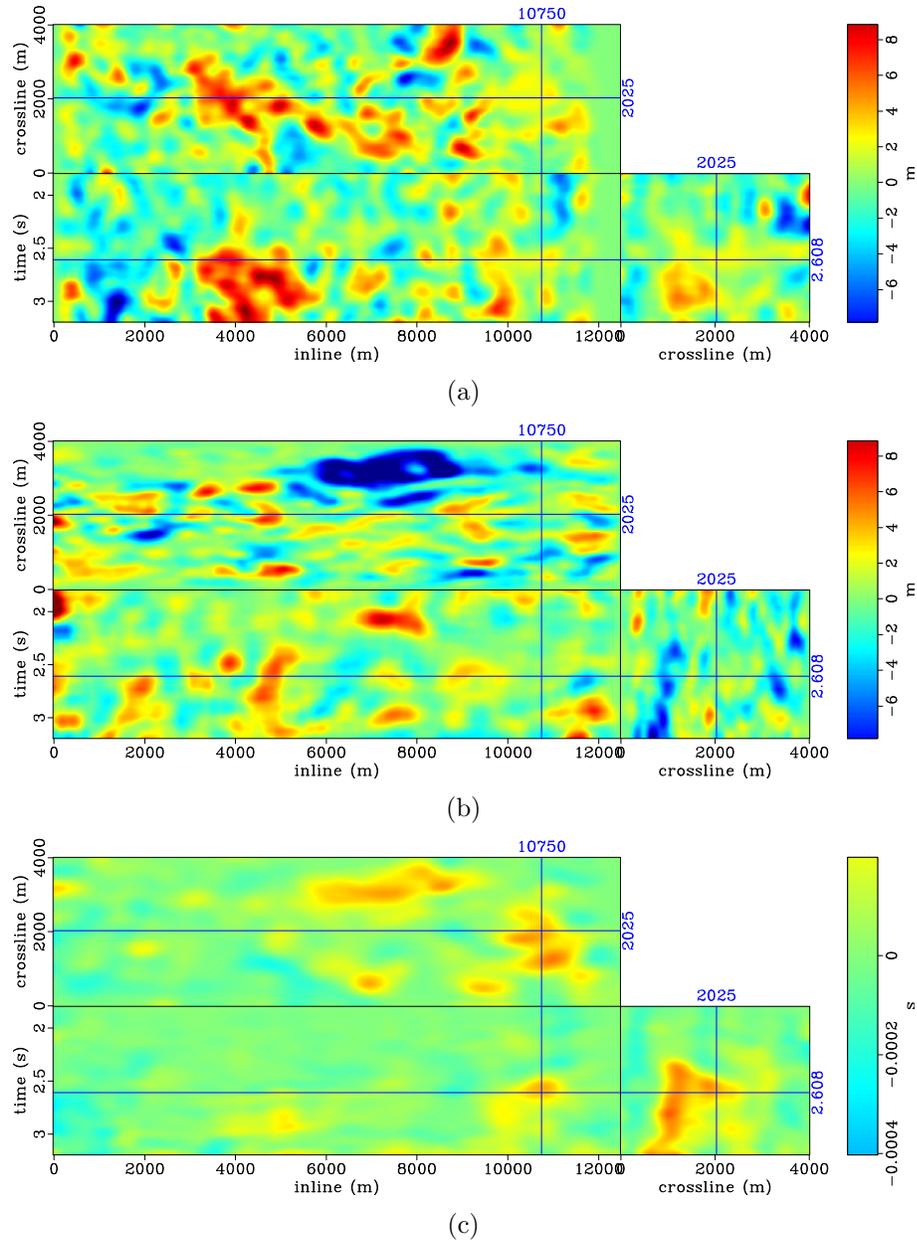


Figure 2.7: Horizontal displacement components between the 2006 monitor and the baseline in the inline (a), and crossline (b) directions; Time-shift errors (c) that are caused by neglecting horizontal displacements. [CR].

chap2/. x-20063,y-20063,ts-d-3

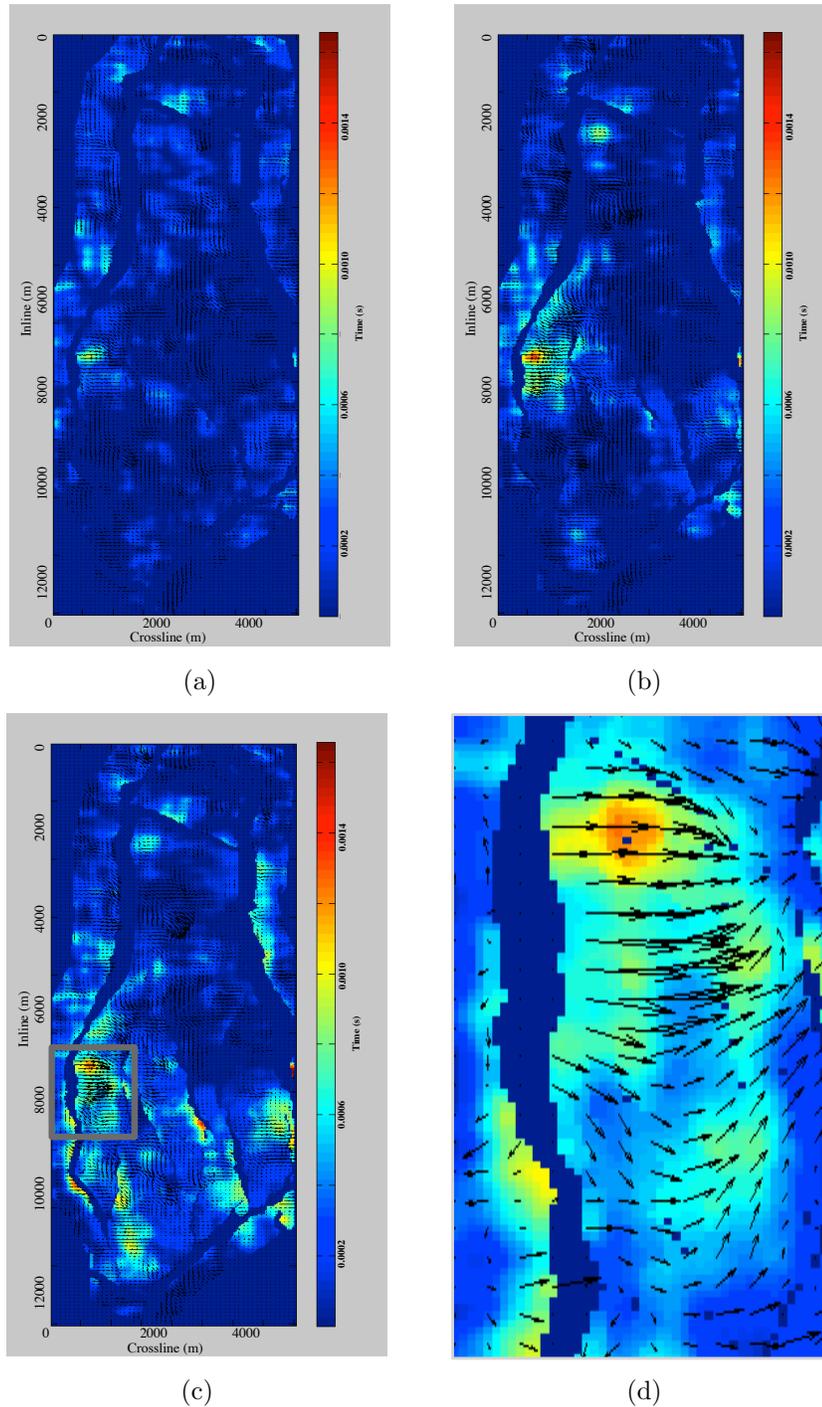


Figure 2.8: Absolute vertical and lateral displacements at the top of the reservoir between the baseline and the 2003 (a), 2004 (b), and 2006 (c) monitor images. In (c), the box indicates the location of the zoomed panel (d). The arrows indicate the displacement direction. [CR]. `chap2/. shift-top1,shift-top2,shift-top3,shift-top4`

Velocity changes between the baseline and the three monitor surveys, computed from equation 2.3, are presented in Figure 2.9. Figure 2.10 shows errors in the estimated velocity change that result from neglecting lateral displacements. Note that as production-induced changes increase, the errors introduced by considering only vertical displacements increase. Figure 2.11 shows the velocity changes between the baseline and the three monitor surveys extracted along the top of the reservoir (Figure 2.9). Velocity errors along the top of the reservoir caused by considering only vertical displacements are shown in Figure 2.12.

Before computing the final amplitude differences between the baseline image and the warped monitor images, I attenuate residual artifacts by the optimized matched filtering method described in a previous section. The starting parameters—derived by trial and error—and the optimization bounds are presented in Table 2.1. Selected optimized parameters derived from matching the baseline image to the 2006 monitor image are shown in Figure 2.13. The filters are estimated within a window that is 1.0 s long above the reservoir and a 0.6 s long below it. Note that the smoothness constraint on the filtering parameters ensures that the parameters (and hence the filters) do not change significantly from trace to trace.

Table 2.1: Matched filter parameters

	Lower bound	Starting guess	Upper bound
Filter length	0.50 s	0.65 s	0.80 s
Upper boundary perturbation	-0.08 s	0.00 s	0.08 s
Lower boundary perturbation	-0.08 s	0.00 s	0.08 s
Damping parameter	1×10^{-4}	1×10^{-3}	1×10^{-2}

Figures 2.14 shows the time-lapse images between the baseline and the 2006 monitor obtained after different processing steps. Note how undesired artifacts in the time-lapse image are attenuated by warping and matched filtering. Note that the large amplitude differences caused by imaging the monitor data with the baseline velocity have been attenuated by warping (Figures 2.14(b)). In addition, note that residual amplitude differences generated by non-repeatability artifacts in the two seismic images have been attenuated by matched filtering (Figures 2.14(c)).

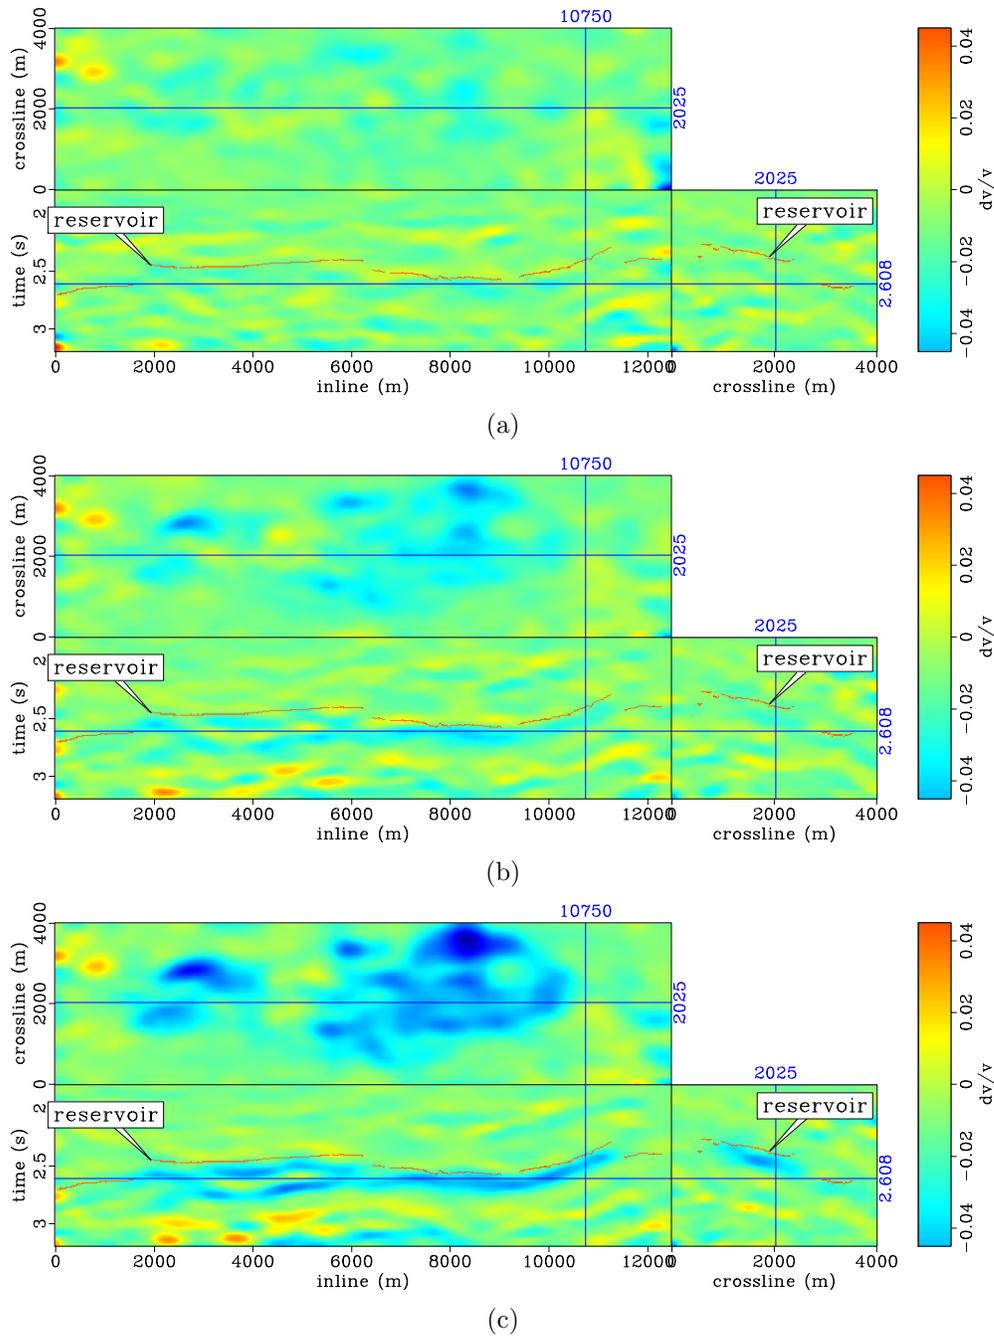


Figure 2.9: Fractional velocity change between the baseline and the 2003 (a), 2004 (b) and 2006 (c) monitor images. Note that in general, due to depletion, velocity within the reservoir decreases with time. The velocity change along the horizon (top of the reservoir) for these three cases are shown in Figure 2.11. [CR].
 chap2/. vv-full-2003-f,vv-full-2004-f,vv-full-2006-f

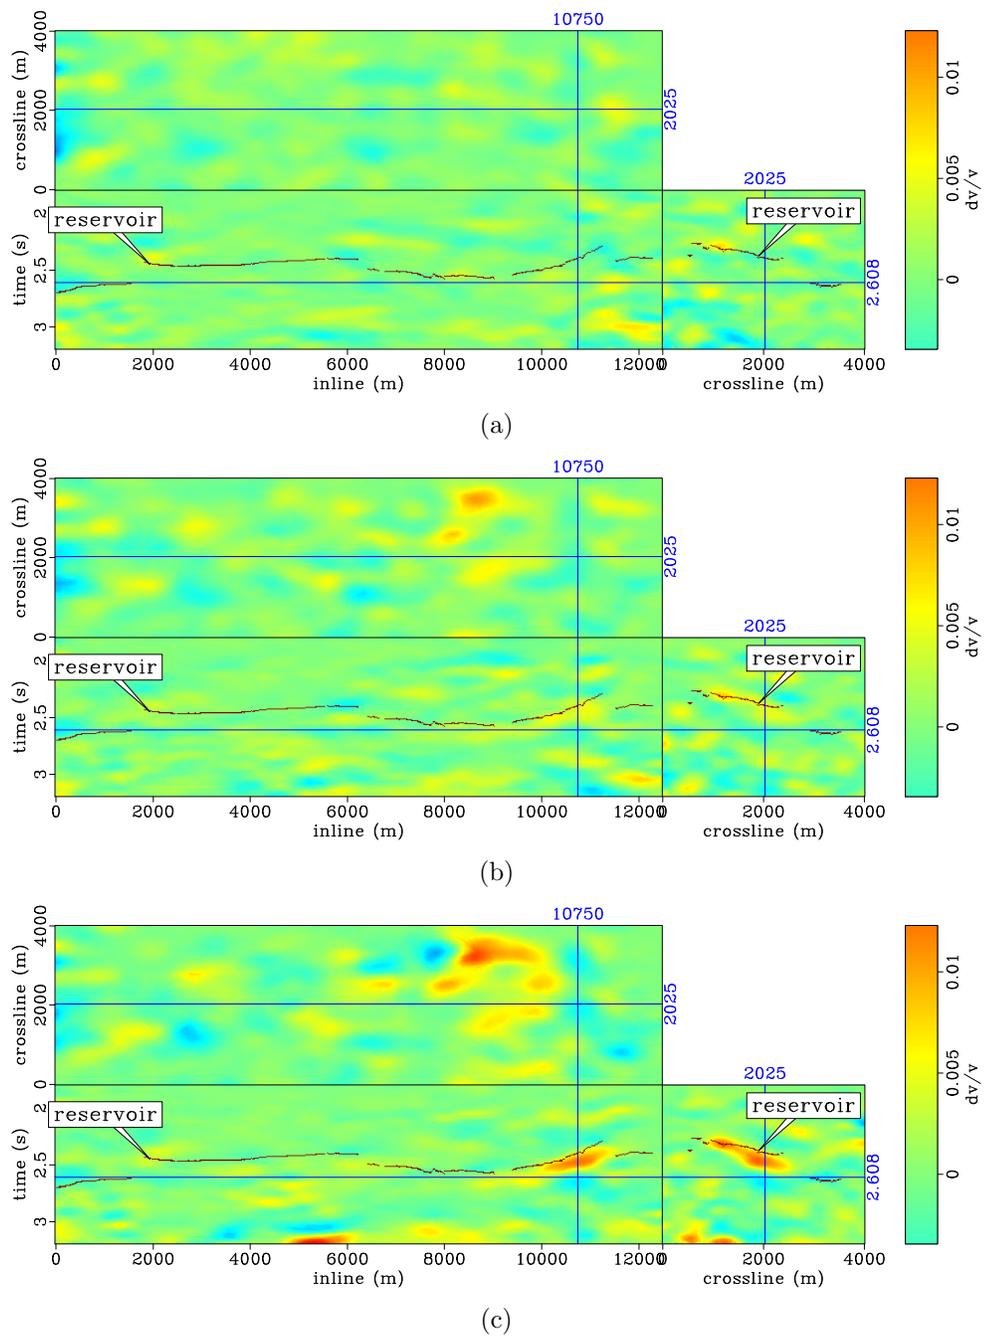


Figure 2.10: Velocity errors caused by neglecting horizontal displacements between the baseline and the 2003 (a), 2004 (b) and 2006 (c) monitor images. Note that the error in velocity increases with time. The errors in velocity change along the horizon (top of the reservoir) are shown in Figure 2.12. [CR].
 chap2/. vv-d-full-2003,vv-d-full-2004,vv-d-full-2006

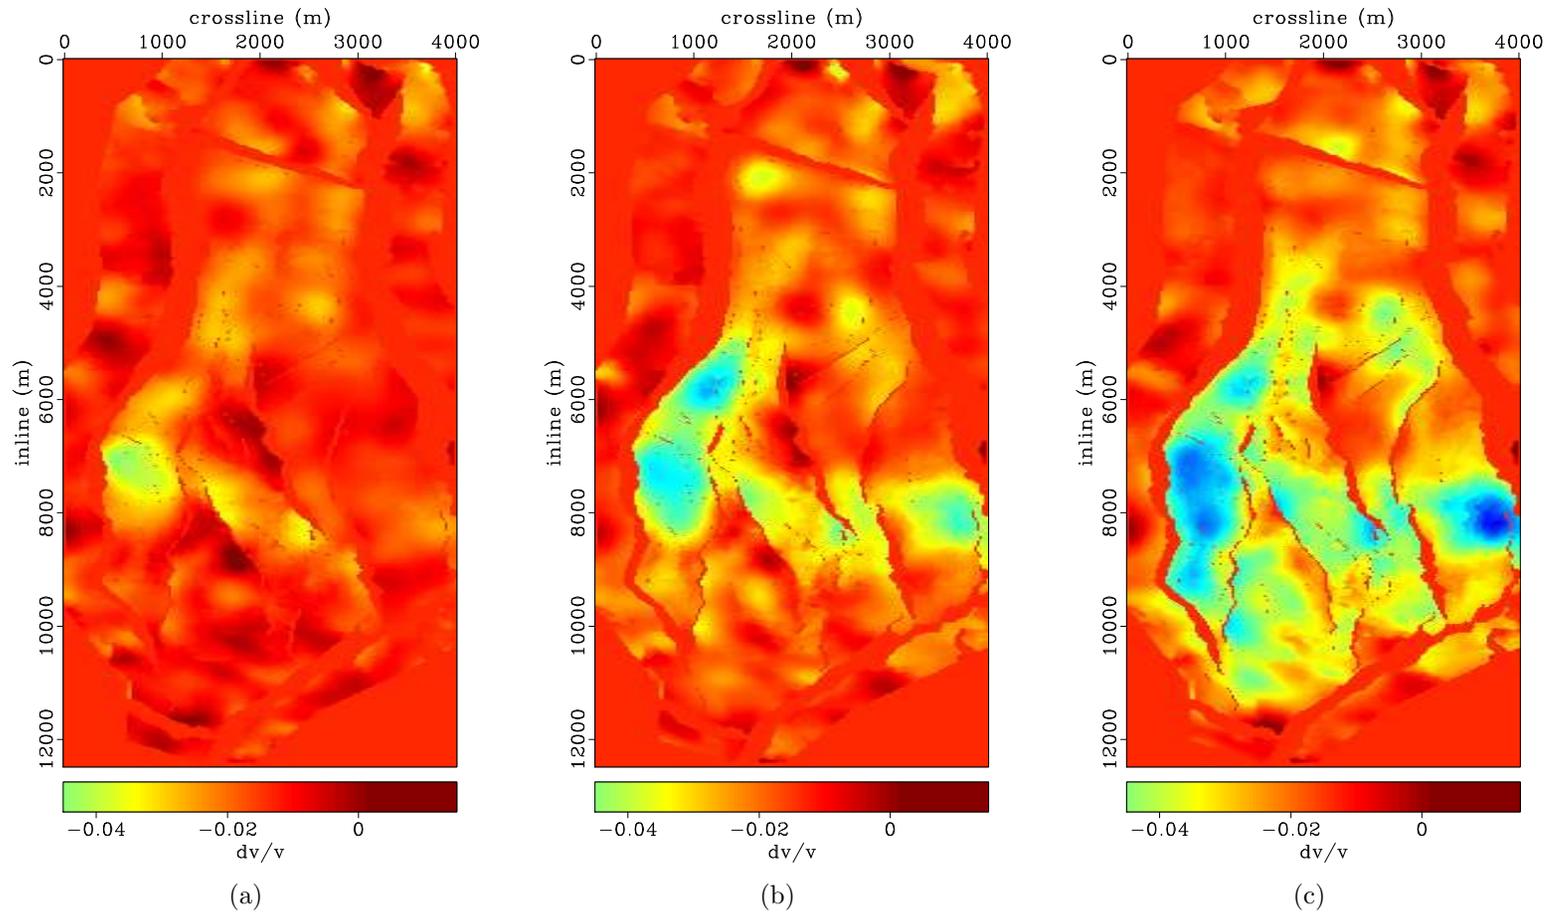


Figure 2.11: Maps of fractional velocity change at the top of the reservoir (Figure 2.9) between the baseline and the 2003 (a), 2004 (b) and 2006 (c) monitor images. Note that in general, due to depletion, velocity within the reservoir decreases with time. [CR]. `chap2/. vv-map-2003,vv-map-2004,vv-map-2006`

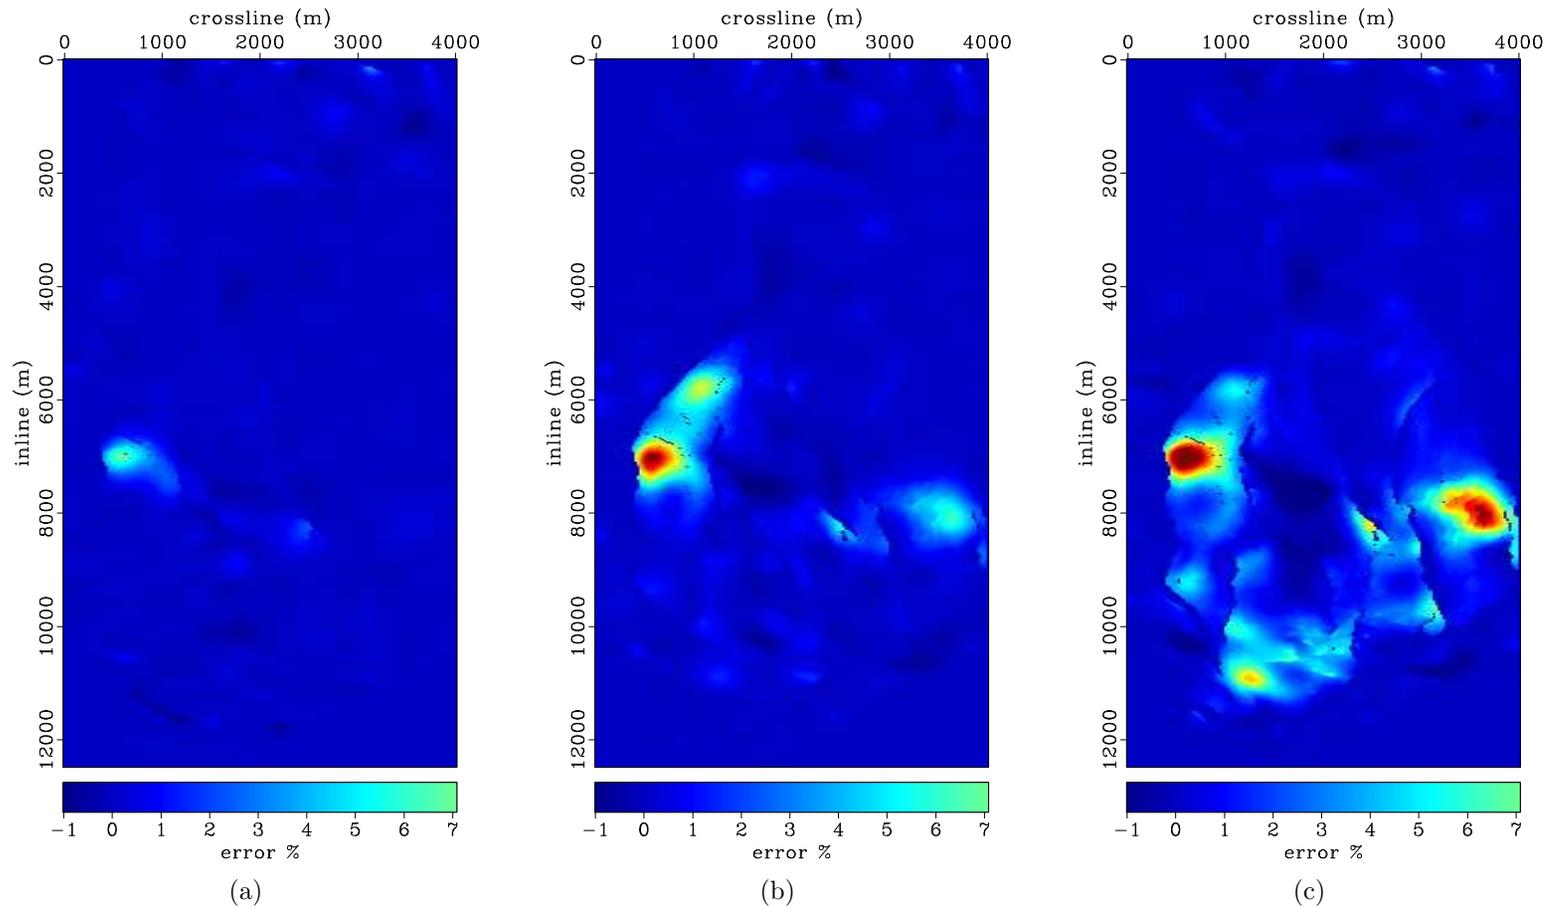


Figure 2.12: Maps of errors in fractional velocity change at the top of the reservoir (Figure 2.10) between the baseline and the 2003 (a), 2004 (b) and 2006 (c) monitor images. Note that the errors are greatest in parts of the field with the largest change (Figure 2.11). Furthermore, note how the errors increase with time. [CR].

chap2/. vv-d-map-2003,vv-d-map-2004,vv-d-map-2006

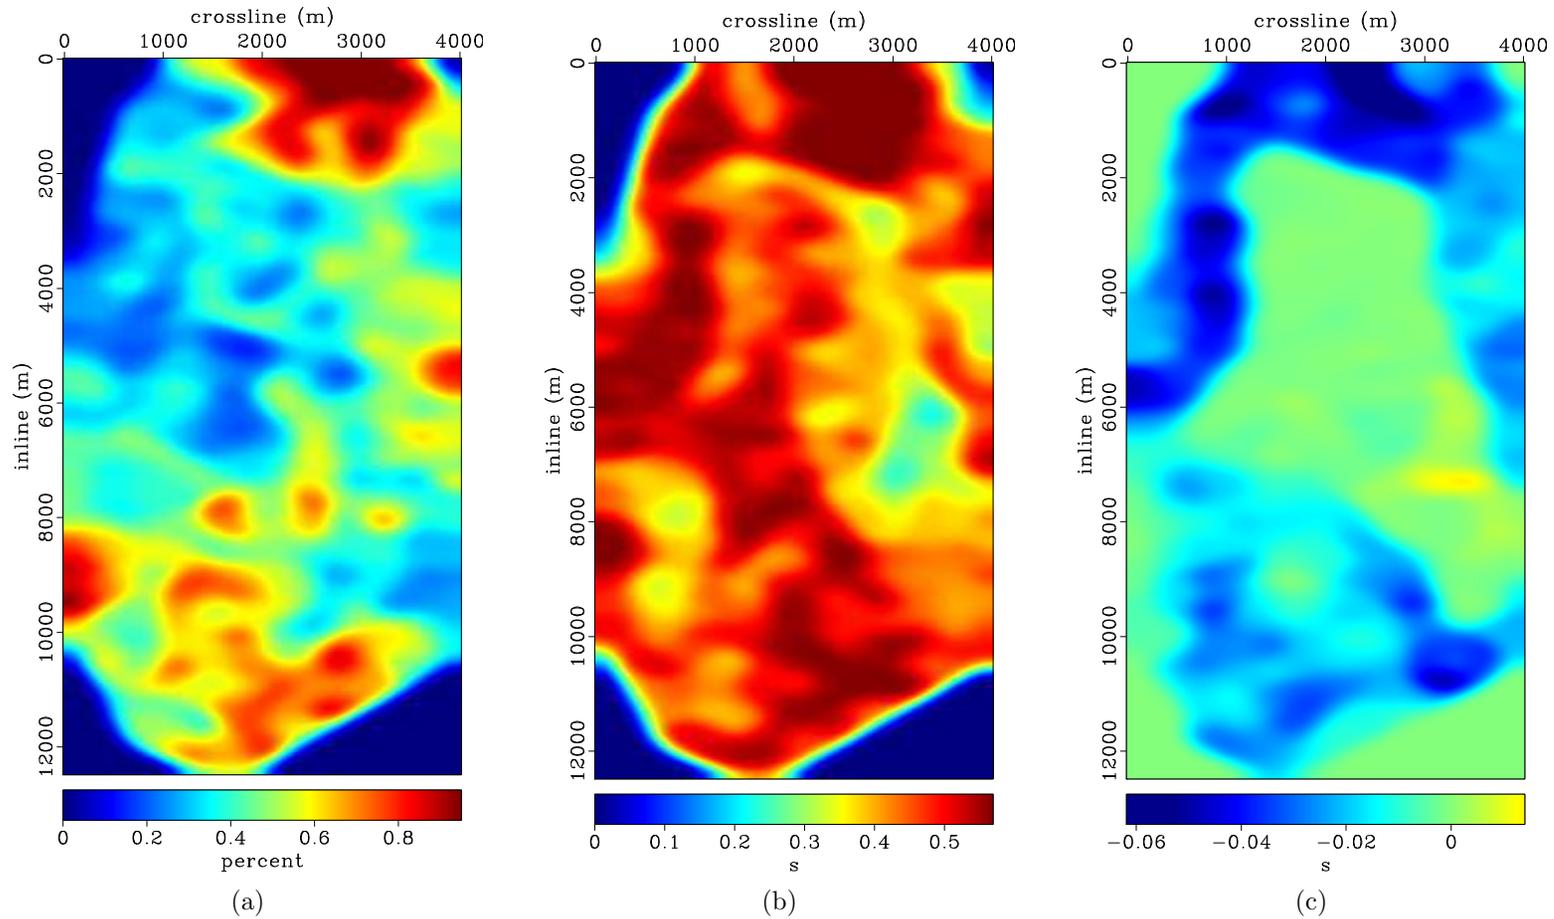
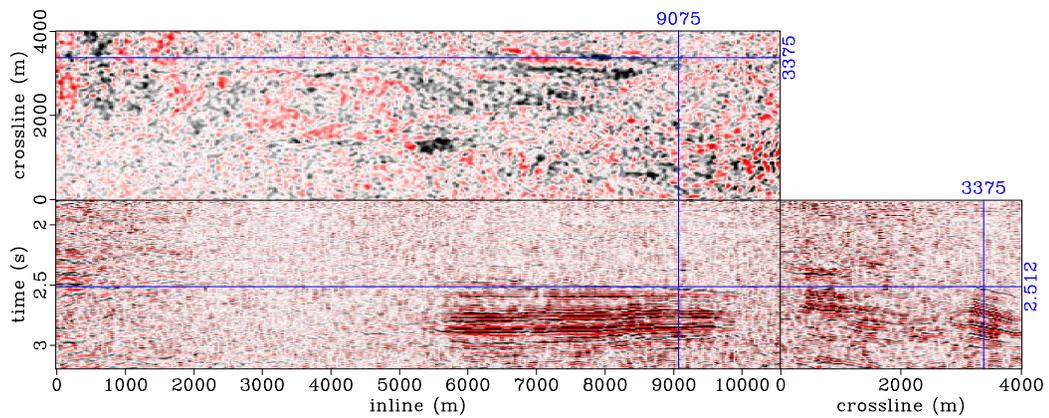
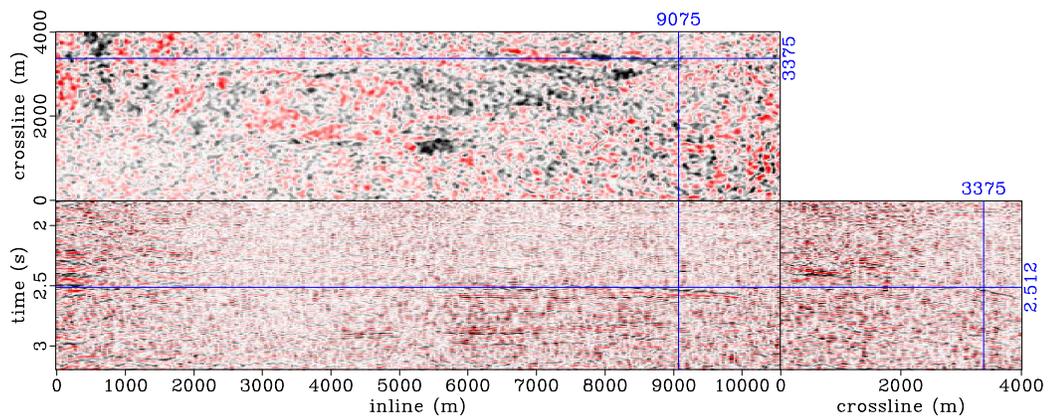


Figure 2.13: Maps of selected optimized matched filtering parameters. The panels show the damping parameter ϵ (a), the filter length, and perturbations to the upper boundary of the estimation window. [CR].

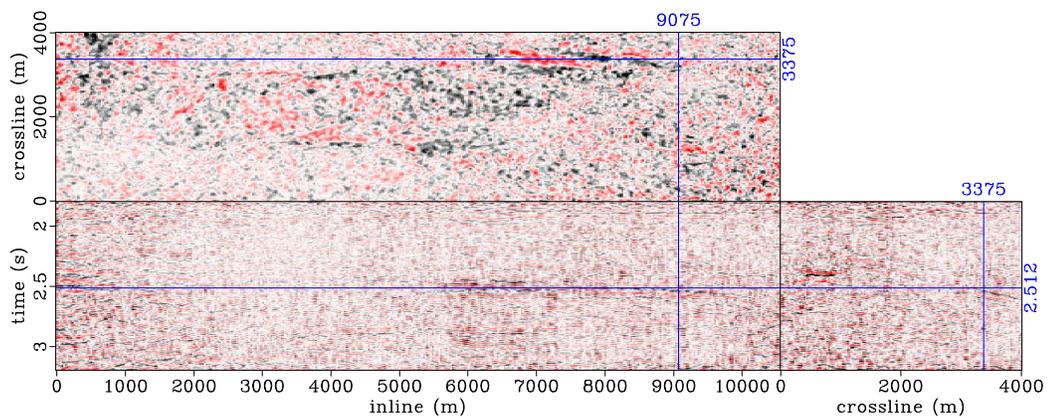
chap2/. map-eps,map-filt,map-nup



(a)



(b)



(c)

Figure 2.14: Time-lapse images between the 2006 monitor and the baseline after different processing steps. The panels show the time-lapse image before warping (a), after warping (b), and after optimized matched filtering. The time-lapse amplitudes of interest are located around the time-slice position. Note the improvements in quality of the time-lapse image after different processing steps. [CR].

chap2/. 4d-or0-2006,4d-org-2006,4d-opt-2006

Even when filter parameters have been selected carefully, using a single set of parameters will not provide optimal filters at all locations. For example, Figure 2.15 shows comparisons among enlarged sections of the time-lapse images after warping, after matched filtering with a single set of parameters, and after matched filtering with optimized parameters. In Figure 2.15(b), the reservoir time-lapse amplitudes are contaminated by vertical streak of filtering artifacts. Artifacts, such as these, are one reason why some practitioners argue avoid matched filtering in seismic cross-equalization (Lumley et al., 2003). In the time-lapse image obtained by optimized filtering, these artifacts are negligible in amplitude (Figure 2.15(c)).

To better understand the impact of matched filtering on reservoir time-lapse amplitudes, let us consider the average time-lapse amplitudes above, below, and within the reservoir. Figure 2.16 shows time-lapse amplitudes in a 0.6 s window around the reservoir. To obtain these (and similar) maps of average time-lapse amplitudes, I compute the difference between the sums of the absolute amplitudes in baseline and monitor images—within the specified window. The time-lapse amplitude maps in Figure 2.16 contain production-related changes overlain by varying levels of contaminating non-repeatability artifacts. In this figure, the most significant production-related amplitude changes are expected within ovals A and C, and in the segment labeled D. Amplitudes within oval B are predominantly non-repeatability artifacts.

Figure 2.16 clearly show one limitation of matched filters in seismic cross-equalization. Because it assumes that artifacts within the estimation window are stationary in time, where this assumption breaks down, matched filtering generates artifacts that further contaminate production-related differences. Comparing Figures 2.16(b) to 2.16(a), within the oval labeled A, note that matched-filtering leads to a improbable large increase in the time-lapse amplitudes. As would be shown later, this increase in time-lapse amplitudes is caused mainly by artifacts resulting from application of matched filters derived from a single set of parameters. Although the time-lapse amplitude map derived from the optimally filtered image contain some artifacts (e.g., within oval B in Figure 2.16(c)), these are fewer and are at relatively lower amplitudes compared to the single-parameter-filtering result (Figure 2.16(b)).

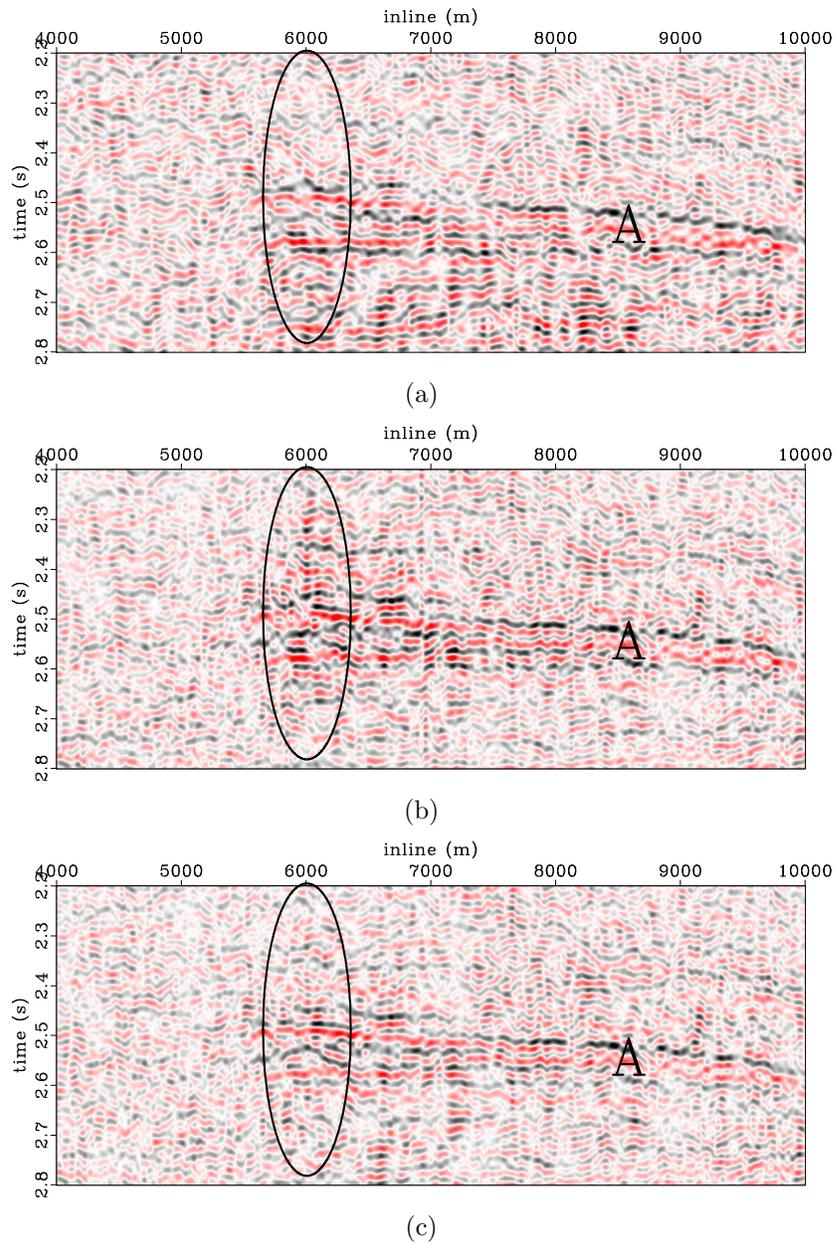


Figure 2.15: Comparison between time-lapse images after different processing steps. The panels show inline sections through the time-lapse image after warping (a), after matched filtering with a single set of parameters (b), and after matched filtering with optimized parameters (c). The time-lapse amplitudes of interest are indicated by label A. In the oval, note that whereas the time-lapse image from filtering with a single set of parameters (b) contains artifacts that are not present in the original time-lapse image (a), these artifacts are attenuated by filtering with optimized parameters (c).

[CR]. chap2/. 4d-org-2006-zm1,4d-smp-2006-zm1,4d-opt-2006-zm1

Figure 2.17 shows the average time-lapse amplitudes within a 0.5 s window above the reservoir. Because no production-related amplitude differences are expected in this non-reservoir region, all amplitudes in the images are undesirable artifacts. Note that the high amplitude artifacts within the oval labeled B in Figure 2.17(a) have been attenuated by matched filtering with a single set of parameters (Figure 2.17(b)). Furthermore, note that the optimally filtered results contain fewer artifacts in this part of the image (Figure 2.17(c)).

Although, in general, matched filtering with a single set of parameters attenuates contaminating artifacts in the non-reservoir region, it also introduces new artifacts in other parts. For example, within the oval labeled A in Figure 2.17(b), which is located directly above a producing segment in the field, note that some amplitudes that are not present in unfiltered image (Figure 2.17(a)) have been introduced. Such artifacts will also contaminate the measure time-lapse amplitudes within the reservoir. Note that within the same region in optimally filtered results (Figure 2.17(c)), these artifacts are negligible. In addition, note that in other areas identified in Figure 2.17, optimized matched filtering provides better results than matched filtering with single set of parameters.

Figure 2.18 shows time-lapse amplitudes within a 0.5 s window below the reservoir. As in Figure 2.17, the large amplitude differences within the oval labeled B, and in other parts of Figure 2.18(a) are undesired artifacts. Note that matched filtering with a single set of parameters attenuate these artifacts (Figure 2.18(b)). However, note that in parts of the field labeled A to D, the amplitude map obtained by optimized matched filtering contain fewer artifacts (Figure 2.18(b)).

As described above, if only vertical displacements are considered, time-lapse images will not be properly aligned. Figure 2.19 shows the errors in time-lapse amplitudes that will result if only vertical displacements are considered. Note that, as observed in estimates of velocity change (Figure 2.10), these errors increase as the production-induced differences between surveys increase.

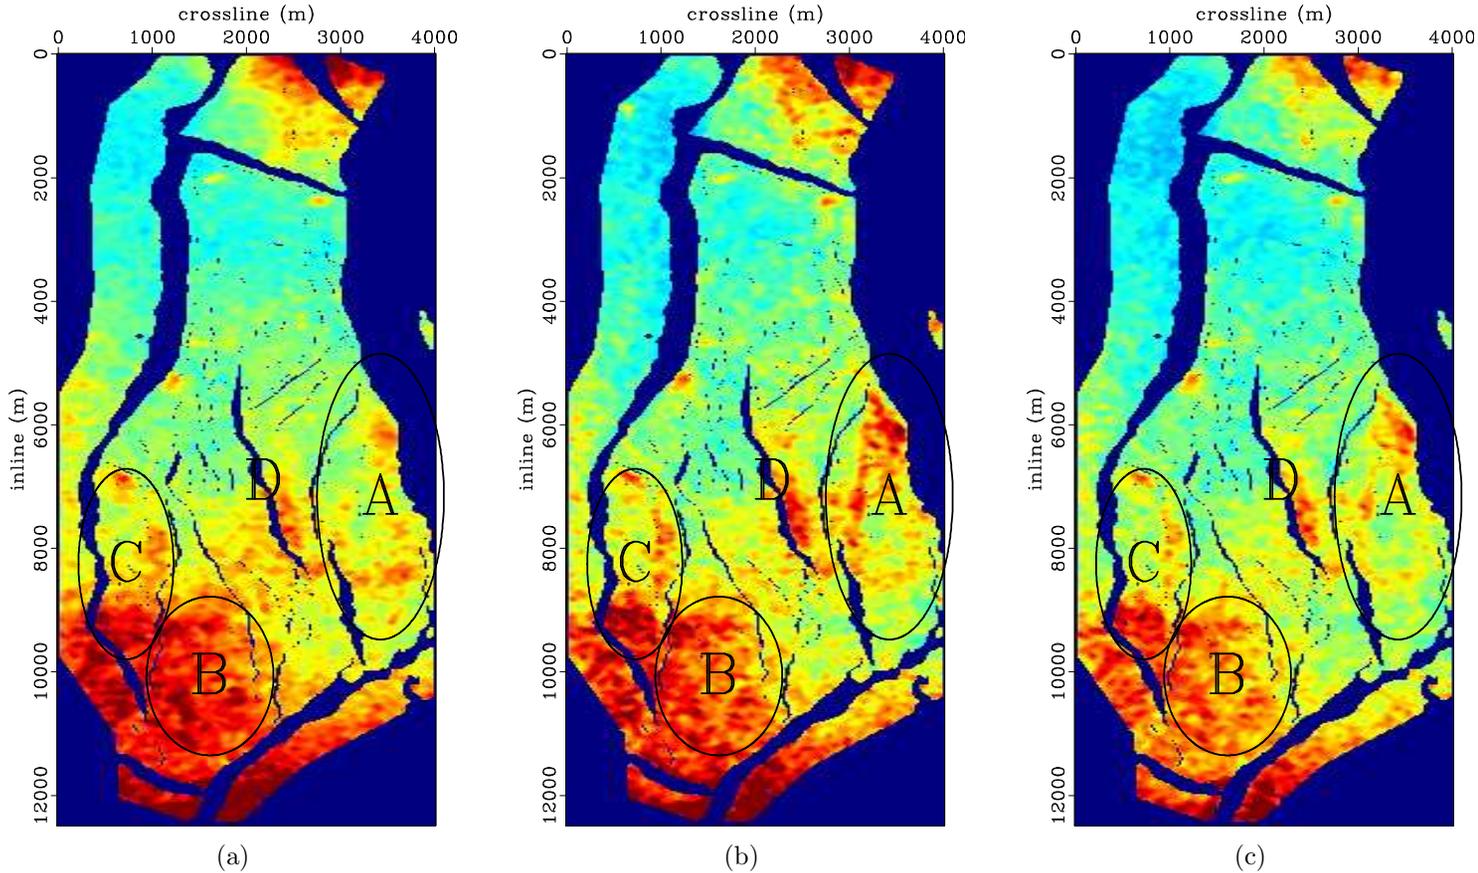


Figure 2.16: Maps of average time-lapse amplitudes between the baseline image and the 2006 monitor image within a 0.6 s window around the reservoir. The panels show amplitude difference after warping (a), after matched filtering with a single set of parameters (b), and after optimized matched filtering (c). As shown in Figures 2.17 and 2.18, amplitudes within oval B are artifacts. The most significant production-related changes are expected only within ovals A and C. These artifacts are better attenuated by optimized filtering (c). Note that the time-lapse amplitudes within oval A in (b) are corrupted by filtering artifacts. [CR]. `chap2/. map-dif-org,map-dif-smp,map-dif-opt`

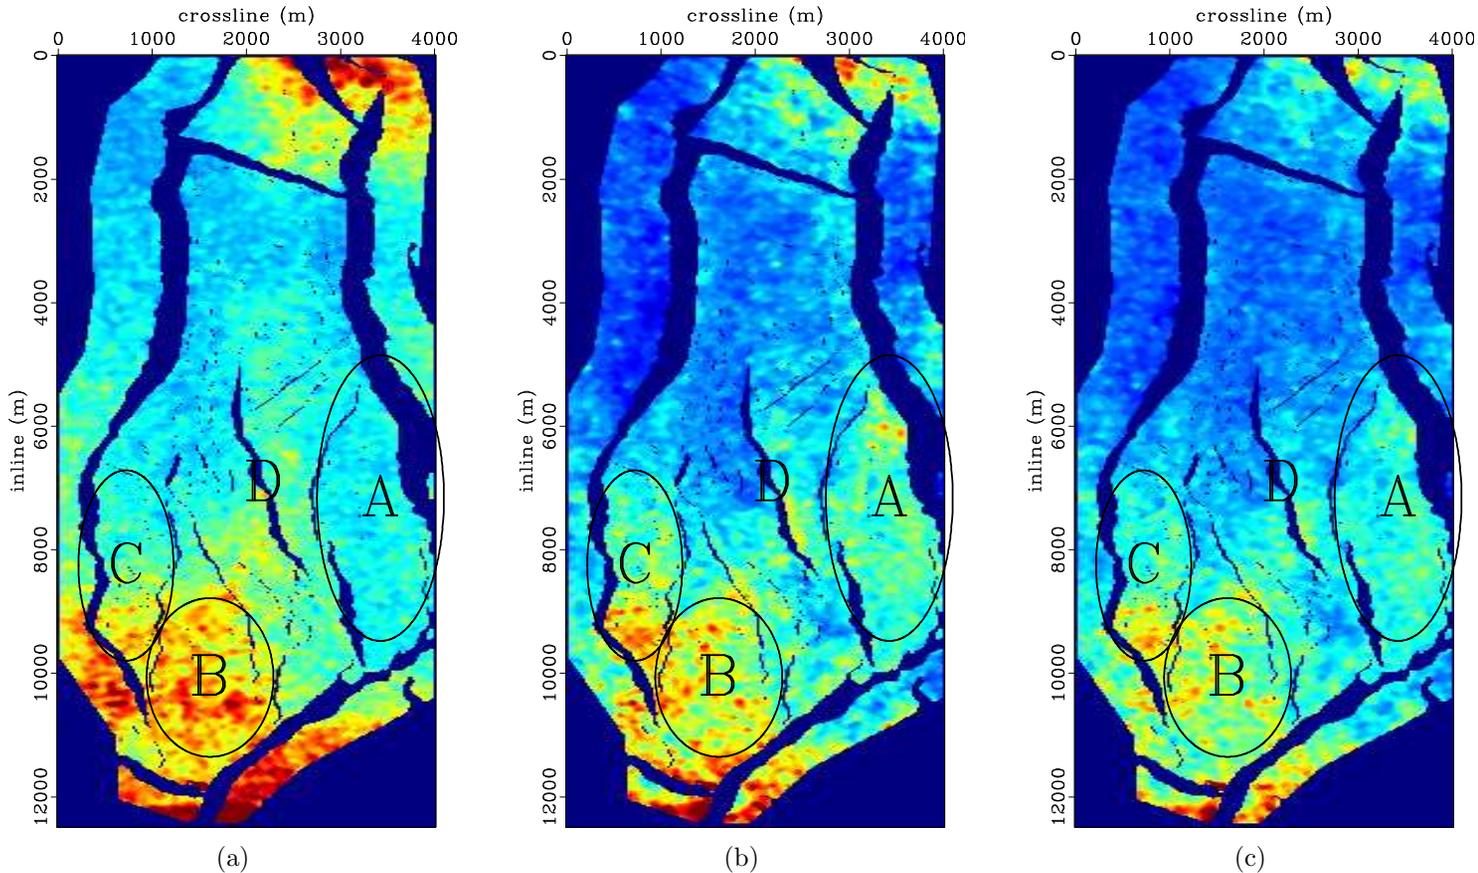


Figure 2.17: Maps of average time-lapse amplitudes between the baseline image and the 2006 monitor image within a 0.5 s window located above the reservoir, where no production-related amplitude change is expected. The panels show amplitude differences after warping (a), after matched filtering with a single set of parameters (b), and after optimized matched filtering (c). In (b), note that matched filtering with a single set of parameters attenuate many of the non-repeatability artifacts (e.g., oval B) but also introduces undesirable artifacts (e.g., oval A). Note in (c) that optimized matched filtering provides satisfactory attenuation of the artifacts without introducing new undesirable artifacts. [CR]. `chap2/. map-dif-org-up,map-dif-smp-up,map-dif-opt-up`

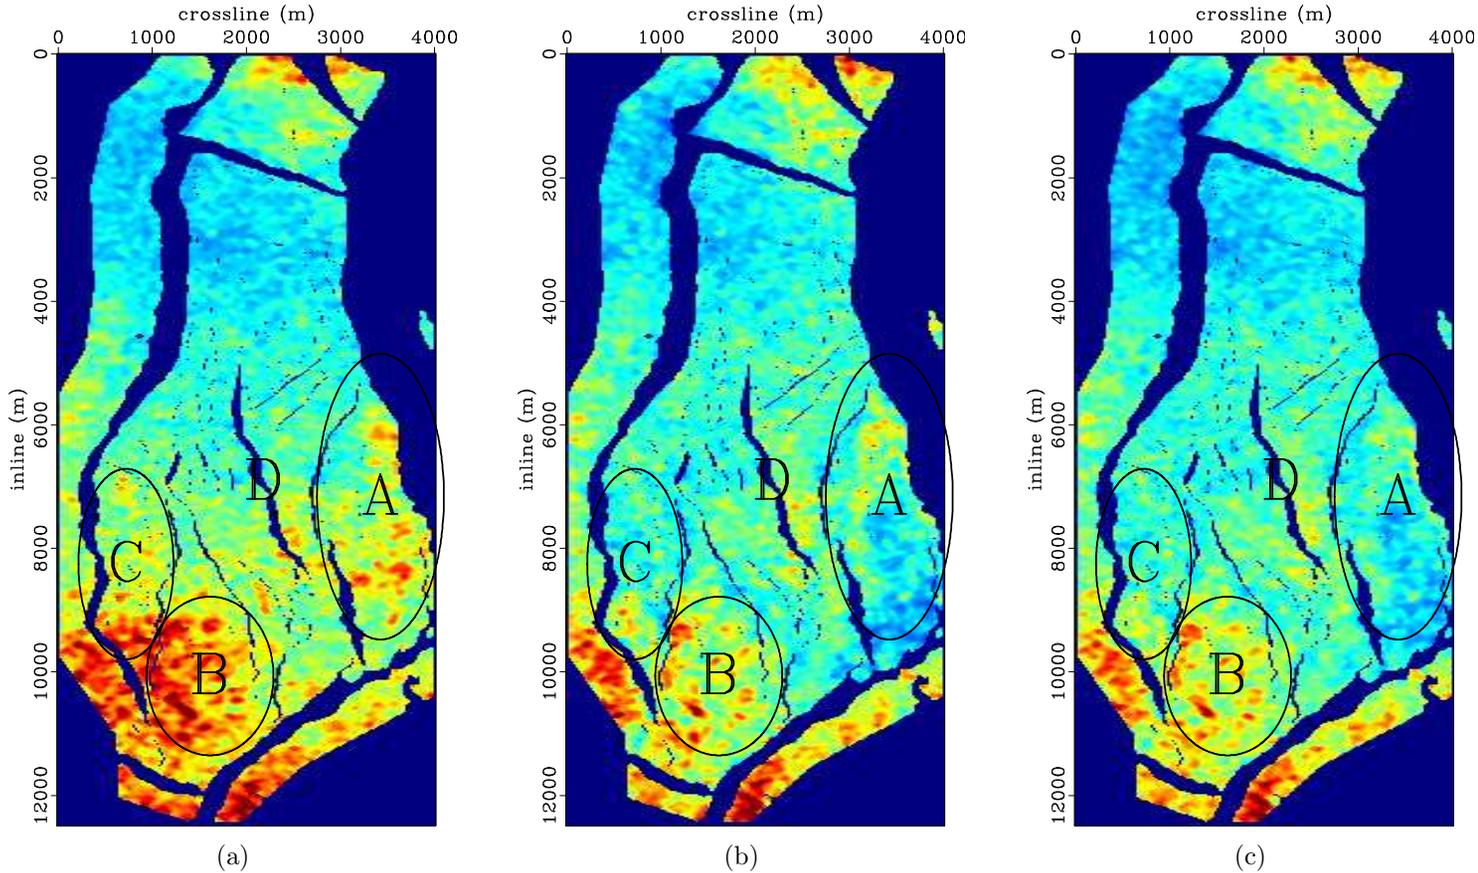


Figure 2.18: Maps of average time-lapse amplitudes between the baseline image and the 2006 monitor image within a 0.5 s window located below the reservoir, where no production-related amplitude change is expected. The panels show amplitude differences after warping (a), after matched filtering with a single set of parameters (b), and after optimized matched filtering (c). Note that matched filtering with provides satisfactory attenuation of the high amplitude artifacts in the (e.g., oval B). Furthermore, note that at several locations (e.g., oval A) the result from optimized filtering contain fewer artifacts. [CR]. chap2/. map-dif-org-dn,map-dif-smp-dn,map-dif-opt-dn

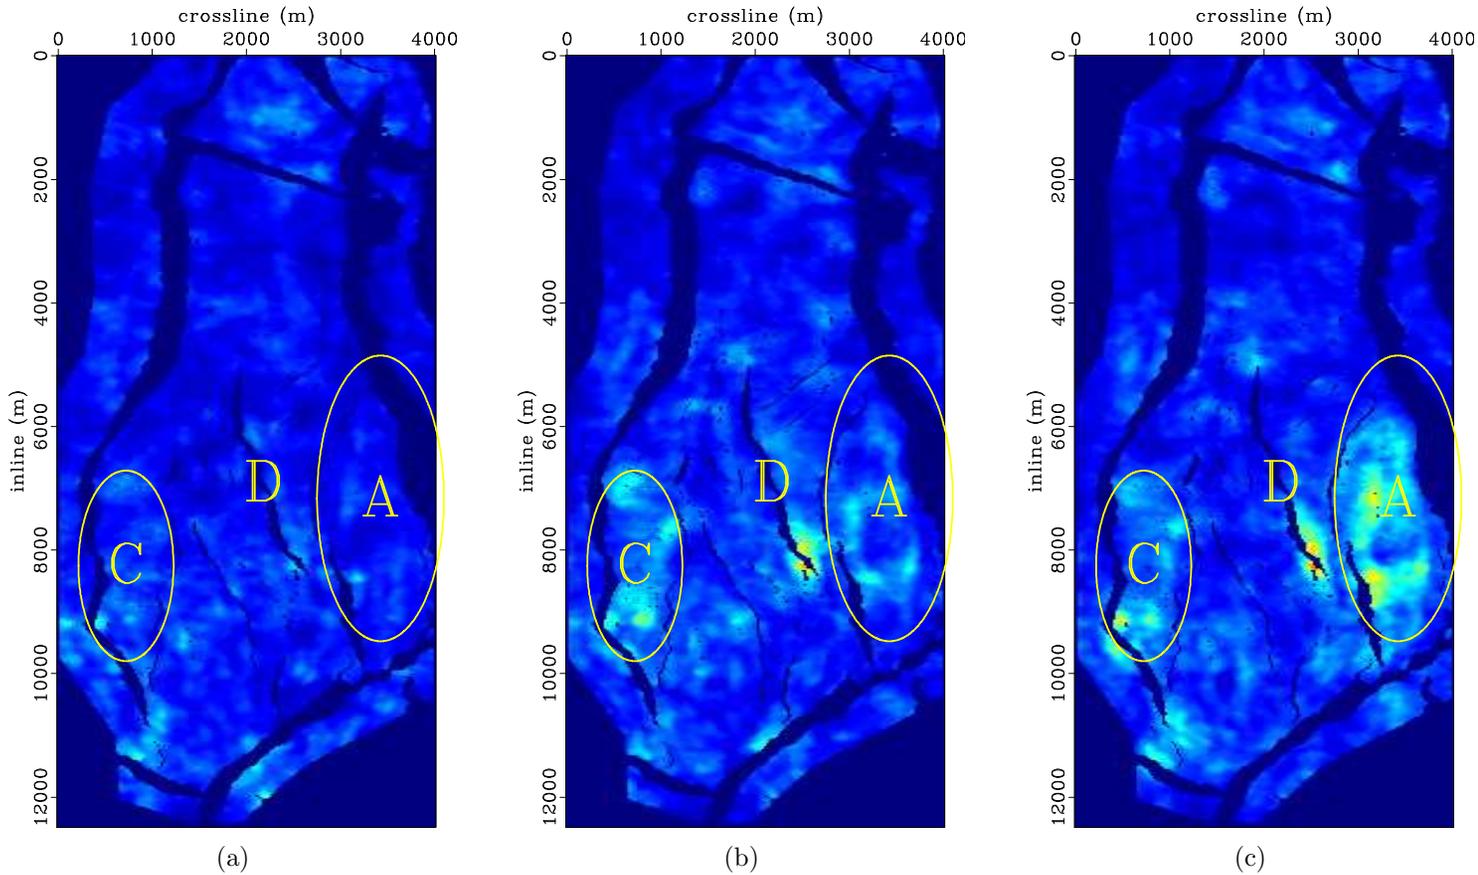


Figure 2.19: Maps showing errors in time-lapse amplitudes within the reservoir caused by considering only vertical displacements versus considering multidimensional displacements. The panels show amplitude differences between the baseline and the 2003 (a), 2004 (b), and 2006 (c) monitor surveys. These amplitude differences represent errors in time-lapse amplitudes caused by approximating multidimensional warping by vertical warping. [CR].

chap2/. map-dif-err1,map-dif-err2,map-dif-err3

DISCUSSION

The sequential one-dimensional warping method described in this chapter provides a robust and efficient alternative to a full, computationally expensive 3D method. By iteratively estimating components of the displacement vectors in a sequential manner using variable correlation gates and lags, it is possible to obtain high-resolution estimates of apparent displacements between time-lapse seismic images (Figures 2.1 to 2.4; Figure 2.6). These apparent displacements contain information about changes within and around the reservoir, which may be interpreted qualitatively (Figures 2.8 and 2.9) or used to build robust geomechanical models (Hatchell and Bourne, 2005a).

Approximating multidimensional apparent displacements between time-lapse seismic images using only vertical time-shifts lead to errors in time-lapse images and their derivatives. In field data example, between 2001 and 2006, lateral displacements of up to 9 m are recorded (Figure 2.7). As shown in Figure 2.7(c), for the field data example, neglecting lateral displacements leads to time-shift overestimation below the reservoir, and time-shift underestimation above the reservoir. Analyses of the estimated displacements show that between 2001 and 2006, velocity within the reservoir has decreased by up to 5% (Figures 2.9 and 2.11). Errors in time-shift, which are more pronounced around dipping reflectors (e.g., Figure 2.7(c)), cause errors of up to 10% in estimates of velocity change (Figures 2.10 and 2.12). Such errors will also propagate to any other reservoir properties (e.g., geomechanical changes) derived from such time-shifts and to time-lapse amplitudes derived by subtracting the warped images (Figure 2.19).

Matched filtering attenuates residual differences that contaminate production-related amplitudes in time-lapse images (Figure 2.14). Although matched filtering with a single parameter set provides a good satisfactory results in parts of the time-lapse image, it also introduces artifacts in other parts (Figure 2.15(b)). However, optimized parameters derived from the proposed method enable computation of the optimal matched filters at each trace location, thereby improving the reliability

of time-lapse amplitudes within the reservoir (Figure 2.15(b)). As shown in Figures 2.17 and 2.18, optimally filtered time-lapse images contain fewer artifacts above and below the reservoir. Therefore, optimized matched filtering provides improved confidence in the time-lapse amplitudes within the reservoir (Figure 2.16).

CONCLUSIONS

In this chapter, I discussed two common cross-equalization methods and how they can be improved. Using the Norne field time-lapse data sets, I showed how these improved methods can provide reliable information about changes within the reservoir.

The sequential one-dimensional warping method, adapted from Hale (2009), allows efficient multidimensional warping of time-lapse seismic images while overcoming cost limitations of full multidimensional warping and avoiding errors associated with vertical warping. Where the estimated apparent displacements are used to derive changes in reservoir properties, and where the reservoir and overburden formations are not horizontal, considering only vertical displacements may cause serious estimation errors. Importantly, this warping method will be used as a preprocessing step for the inversion methods developed in later chapters.

After warping, residual artifacts in time-lapse seismic images caused by non-repeatability between surveys are attenuated by matched filtering with optimized parameters. In this chapter, I showed how optimal filtering parameters can be obtained using an evolutionary programming algorithm. These filters provide more reliable time-lapse images than conventional matched filtering with single parameters. Together, sequential one-dimensional warping and optimized matched filtering form an efficient and robust cross-equalization scheme.

Although the methods described in this chapter perform satisfactorily in conventional time-lapse imaging, such as in the one discussed in this chapter, they are insufficient in many cases. In the next chapter, I discuss joint least-squares inversion methods that overcome some of the limitations of these conventional time-lapse cross-equalization methods.

ACKNOWLEDGEMENTS

I thank Statoil and partners for releasing the Norne data through the NTNU/IO CENTRE. Bard Osdal and Trine Alsos (Statoil), Mohsen Dadashpour (NTNU) and David Echeverria (formerly of Stanford Smart Fields, now IBM) provided assistance throughout this study. Mosab Nasser (formerly of Shell, currently of Maersk Oil), my supervisor during an internship at Shell, suggested that I explore ways to improve local match-filtering of time-lapse data sets.