Extending the useable bandwidth of seismic data with tensor-guided, frequency-dependent filtering

Edward Jenner[1*], Lisa Sanford[2], Hans Ecke[1] and Bruce Golob[1] describe a frequency-dependent filtering technique that can significantly increase the available bandwidth of the seismic data.

Improving the bandwidth of seismic data has been an ongoing endeavour in increasing the usability of seismic data in many plays throughout the world. While various broadband acquisition techniques have been recently developed for marine acquisition, the most significant issue for land seismic has been the low signal-to-noise ratio, particularly at the low and high ends of the frequency spectrum. In this paper, we describe a frequency-dependent filtering technique that can significantly increase the available bandwidth of the seismic data. We show a field data example and demonstrate that the enhanced bandwidth seismic data ties with well logs in the area.

A variety of methods have been proposed to increase the bandwidth of seismic data that fall broadly into three categories:

1. Inversion-type approaches where some pre-determined information or assumption is used to overcome the inherent non-uniqueness of the inversion solution (Zhang and Castangna, 2011).
2. Methods that use the available bandwidth with high signal-to-noise to extrapolate or 'predict' the low and high-frequency components (Smith et al., 2008).
3. Methods that apply some sort of frequency-dependent filtering to the seismic data to improve the signal-to-noise of the high and/or low-frequency components (Whitcome and Hodgson, 2007).

While the third type of frequency enhancement may be seen as relatively lower impact than the inversion or extrapolation schemes, it does not make specific assumptions about the geology or the relationship between low and high frequencies in the data and is thus more objective than the other two methods. In addition, it can form the basis of pre-conditioning for the first two methods potentially increasing their effectiveness if their assumptions are not violated.

The biggest issue with frequency-dependent filtering usually occurs when applying the filter to the very low and high frequencies where the signal-to-noise is so low that the filtering methods cannot separate signal from noise. In this case, we cannot simply apply a filter such as FXY, FK or median filtering with frequency-dependent parameters. However, in recent years several authors have shown the benefits of using structure-oriented filtering in reducing noise while maintaining structural features (Chopra and Marfurt, 2008; Helmore et al., 2007).

Certainly structure-oriented filtering can provide many benefits in terms of noise attenuation. However, in the cases where at some frequencies we have such a low signal-to-noise ratio that we can hardly identify the signal, it has a significant advantage. We can use the structure term from the frequency component with high signal-to-noise to guide the filtering of the frequency components with much lower signal-to-noise. The only assumption here is that the high and low frequencies approximately follow the structure of the stack used to compute the structure tensors. As we shall see, however, there does not have to be a 1:1 correspondence – i.e., we can still extract high-frequency details from structure tensors that are computed from a low-frequency stack. However, the tensors cannot deal with cross-dipping events, so applying this method to data before time or depth migration is not practical.

Method

While various methods that are somewhat similar exist to apply structure-oriented filtering (e.g., Fehmers and Hocker, 2003; Hoeber et al., 2006), we employ the method of Hale (2009a, 2009b). However, we have no reason to suspect that other methods should not be equally as effective. For each sample point, the structure tensor is computed as the smoothed outer products of local image gradients. The eigenvalues of these matrices give us a measure of isotropy, linearity, and planarity at that location, as well as the direction of the identified linear or planar features. The filtering process is computed by a modified non-linear anisotropic diffusion equation, which is somewhat comparable to anisotropic adaptive filtering with the anisotropy defined by the structure tensors.

As illustrated in Figure 1, the general workflow for applying the structure-oriented, frequency-dependent (SOFD)

[1] *NEOS, formerly at ION Geophysical.*
[2] *ION Geophysical.*
[*] *Corresponding author, E-mail: ejenner@neosgeo.com*

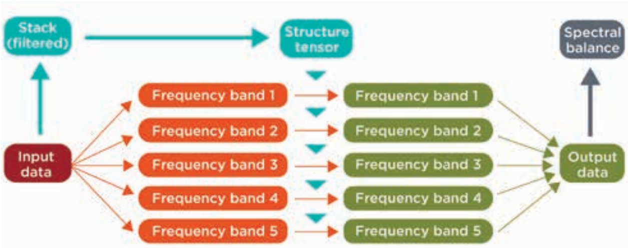

Figure 1 *Schematic of the workflow used to apply the structure-oriented frequency-dependent (SOFD) filtering.*

filtering consists of computing structure tensors from the input data, splitting up the input data into frequency bands, applying frequency-dependent filtering on each frequency band and then recombining the frequency bands to give the output data. The 'stack' used to compute the structure tensors may be produced in any number of ways that gives the desired structure tensors. The filtering process is not particularly sensitive to the stack volume used to compute the tensors, but the stack should be generated to maximize fault/discontinuity definition so that the filtering preserves these discontinuities as much as possible. In some cases it may also be beneficial to bandpass-limit the stack and even use different bandpass filtered stacks for the structure tensors to be used in different frequency bands. In any case, the stack must be of sufficiently high signal-to-noise to accurately compute the structure tensors.

Examples

The data shown are the p-wave data from two ION GeoVentures ResSCAN multi-component surveys, BuffalohornSCAN and GroundhogSCAN. Buffalohorn is a 180 mile^2 survey targeting the Mississippian Lime in Oklahoma and Groundhog is a 400 mile^2 survey targeting the Marcellus in Pennsylvania. An advanced processing sequence including 5D interpolation, anisotropic velocity analysis and azimuthally anisotropic PSTM (Jenner, 2011) was used before the SOFD filtering.

Figure 2 shows an example of applying this process to post-stack data from the Buffalohorn survey, in this case a 3D pre-stack time migrated image. The top panel shows a profile through the input stack (Figure 2a) and the same stack bandpass filtered into four frequency bands; 2-3-8-10 Hz, 8-10-60-70 Hz, 60-70-80-100 Hz and 80-100-130-160 Hz (Figures 2b-2e respectively). These bands were chosen to illustrate the variation in signal-to-noise with frequency and are not the bands used in the filtering process. Frequency band 2 (10-60 Hz) is the main frequency band of the data with good signal-to-noise, frequency bands 1 and 3 have lower signal-to-noise and frequency band 4 has very low signal-to-noise. The frequency panels have been gained for display. However, it is apparent most of the energy observed in the stack is in the 10-60 Hz band. Except for ground roll suppression, most noise attenuation has been performed in this frequency range. Therefore, it makes sense then that the data in this frequency band have a relatively good signal-to-noise.

After applying the SOFD (Figures 2f-2j), the resulting stack is very similar with just a small amount of high frequency noise attenuation apparent. Indeed, very mild filtering was required in the 10-60 Hz band. Stronger filtering was applied in the 2-10 Hz and 60-100 Hz bands and even stronger filtering was applied above 100 Hz.

With the stack now having significantly improved signal-to-noise at the low and high frequencies, processes such as Q compensation, spectral shaping and spectral balance that boost the high and/or low frequencies relative to the mid-frequency band may be employed without generating a large amount of unwanted noise and artifacts in the data. For instance, we applied an amplitude-preserving spectral balance to the output of the SOFD filtering in Figure 2f. This spectral balance preserves the amplitudes in the mid-frequency band such that if the data are bandpass filtered back to the original spectrum, the amplitudes and AVO behaviour will be identical to the input data. In addition, peaks and notches in the spectrum are maintained. Outside the main bandwidth of the data, the character of the spectrum is also maintained, but the amplitudes are boosted such that the overall spectrum of the data is flattened – i.e., very approximately white over a large frequency range. Usually a relatively flat spectrum will

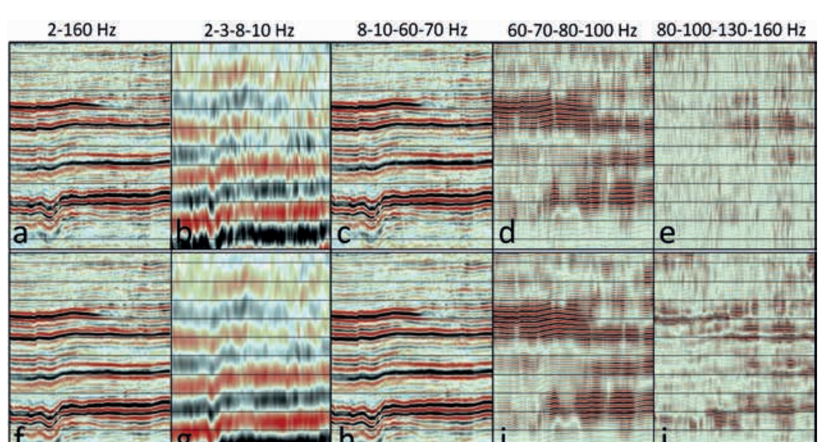

Figure 2 *3D poststack example before (a) and after (f) SOFD filtering. Other panels show bandpass filtered versions before (top row; b-e) and after (bottom row; g-j) SOFD filtering.*

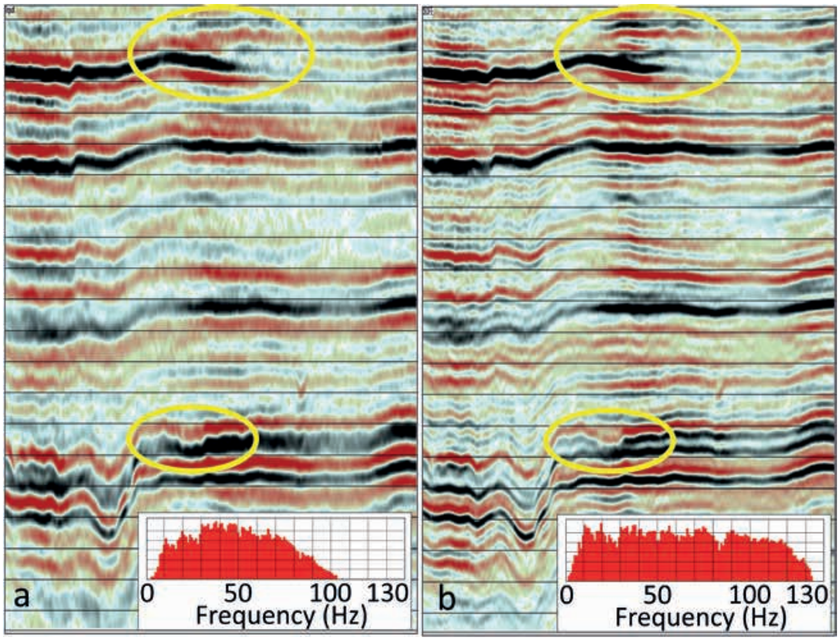

Figure 3 *Applying spectral balance after SOFD filtering. (a) Profile through a 3D stack SOFD filtering. (b) The same data in (a) after amplitude-preserving spectral balance.*

tie the well data and geology. However, if the geology is not approximately white across a wide-frequency band (e.g., it is lacking in frequencies from 60-80 Hz) then the spectrum can be shaped to match the geology from well data.

Figure 3 shows the same profile used in Figure 2 with the SOFD filtering before and after the amplitude-preserving spectral balance. In this example, all the structure tensors were computed from the stack in Figure 3a and applied to the same stack, but in the frequency-dependent manner illustrated in Figure 1. Obviously, the section with the spectral balance has a higher and lower frequency content. However, note that the spectrum of the output matches the spectrum of the input between ~30-60 Hz and that the character (peaks and notches) of the spectrum are preserved outside this frequency band. In addition, boosting the high and low frequencies, which now have a relatively high signal-to-noise ratio, has not introduced significant amounts of unwanted noise and artifacts into the data.

An important point regarding this process can also be made here. Notice in the areas marked with yellow ovals that the high-frequency structure is not identical to the structure in the mid-range frequency band that was used to compute the structure tensors. This is an important point. Although we compute structure tensors on the mid-frequency band and we make the assumption that the low and high frequencies have the same structure, this assumption can clearly be violated to some extent. In effect, the structure tensors 'guide' the filtering, but they cannot create events, nor can they force the high-frequency events to line up exactly with the tensors. In addition, as we shall show below, the tensors cannot create high-frequency events out of random noise in the data. In practice, what this means is that we can only recover high and low frequencies where there is some signal that has approximately the same structure as the computed tensors. However, it does not mean that we will destroy or smear through wedge, on-lap or off-lap stratigraphy that can only be observed in the high-frequency component of the seismic data.

As an illustration of these last points, Figure 4 shows the result of applying this technique to a synthetic section consisting of band-limited random noise and seven horizontal reflectors. The structure tensors were computed using the stack

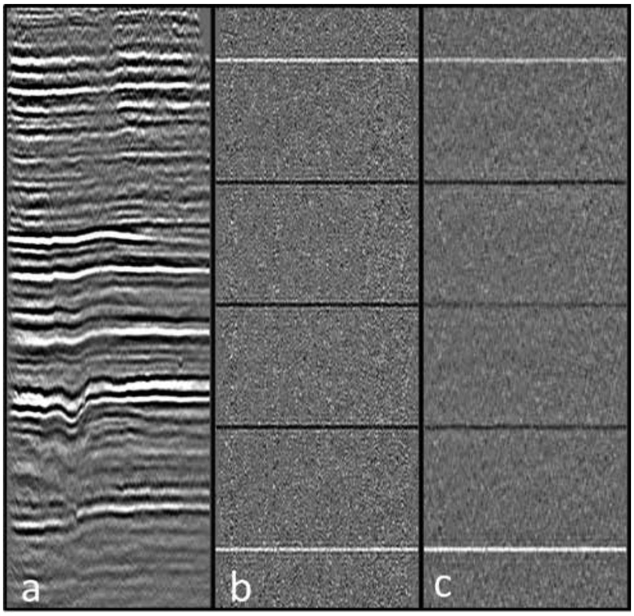

Figure 4 *Applying SOFD filtering to a synthetic with noise. (a) The stack used to compute the structure tensors; (b) the synthetic model; (c) the result of applying SOFD filtering to (b) using the tensors computed from (a).*

in Figure 4a and the SOFD filtering was applied in exactly the same manner to the data in Figure 4b as used in Figure 3, producing the result in Figure 4c. While the random noise is reduced (mostly the high and low-frequency components of the random noise), the horizontal reflectors remain horizontal. In particular, the highly filtered frequency components do not separate from the mid-frequency component despite the tensors not being parallel to the horizontal events. Where there are significant structural differences between the stack used to compute the tensors and the horizontal reflections, the horizontal reflections are reduced in amplitude.

It is important that any bandwidth enhancement technique ties with well data, at least as well as expected from well tie to the input or narrow bandwidth data. Since this method is not performing any extrapolation, prediction or making specific assumptions about the reflectivity, we should expect this to be the case and indeed this is our general experience. As an example, Figure 5 shows a well synthetic before and after the SOFD filtering. For the synthetic tie in Figure 5a, the data were spectrally balanced as far as considered reasonable and then noise attenuation was applied to attenuate high-frequency noise observed after the spectral balance. Clearly, the data tie the well synthetic very well and resulted in a correlation coefficient of 0.69. Figure 5b shows the tie of the data after SOFD filtering and spectral balance but now the spectral balance can recover significantly higher (and some lower) frequencies than before. The correlation coefficient is reduced slightly to 0.63 as is expected whenever correlating a higher frequency trace with a synthetic; however, the tie is still very good. In addition, new events that have been broken out on the seismic data are clearly also evident in the higher frequency well synthetic.

The technique can also be applied to pre-stack data by using the stack to compute the structure tensors and then applying them in a frequency-dependent manner to each migrated offset as a separate volume. This will maintain azimuthal information in the case of migrating into offset-vector tiles (OVT) or azimuth sectors as no filtering is performed across offsets or azimuths. Obviously, additional filtering can be employed either before or after the SOFD filtering as desired based on the requirements for further processing or data analysis.

As an example of applying the technique to prestack data, Figure 6 shows a PSTM OVT-migrated gather from the Groundhog survey with spectral balancing applied before and after the SOFD filtering. The filtering significantly improves the signal-to-noise and usability of the gathers. However, in this example, we purposely show a gather that has some small residual moveout to demonstrate that the gather moveout and characteristics are preserved. In other words, as with the stack example, the fact that the structure tensors computed on the stack do not exactly line up with the pre-stack seismic data does not impede the process. On the other hand, because of the increased resolution, it will often be necessary to perform additional gather conditioning (e.g., residual NMO, alignment) before using the enhanced broadband gathers for inversion or pre-stack attribute analysis.

Stacking the data with and without the SOFD filtering is shown in Figure 7. Since the spectral balancing was applied pre-stack, the effect of stacking the data has reduced the high-frequency noise introduced in the spectral balance process such that the stack does not look particularly noisy. However, the spectra show that the available energy in the stack (Figure 7a) drops off significantly after 75 Hz. In comparison, the stack of the data with SOFD filtering (Figure 7b) has significantly higher frequency content and higher overall signal-to-noise because the frequency components above 75 Hz have now significantly contributed to the stack.

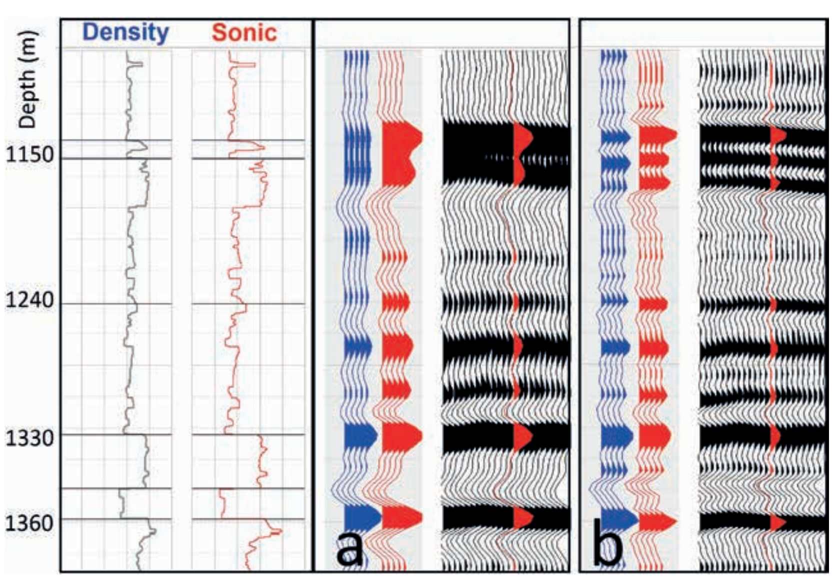

***Figure 5** Well tie of the seismic data to synthetic data computed from density and sonic logs. Blue – synthetic seismic from well logs, red – average of seismic traces in vicinity of well. (a) The best well tie before applying SOFD filtering. (b) The well tie after applying SOFD filtering.*

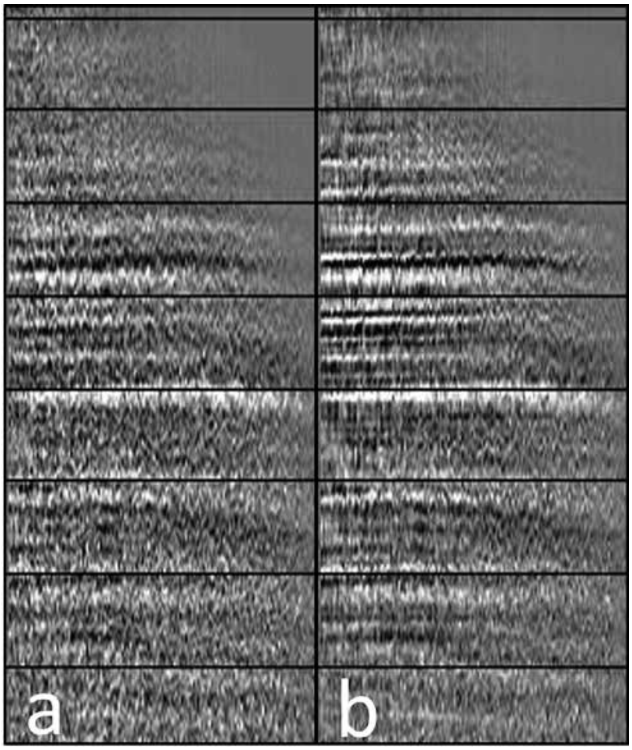

Figure 6 *Common image offset vector tile gather after spectral balance. (a) Without SOFD filtering; (b) with SOFD filtering. In this example, all filtering was applied in the surface X-Y domain with no filtering applied across offsets.*

Finally, we show an example of performing pre-stack inversion of the Buffalohorn survey after applying the SOFD filtering followed by amplitude-preserving spectral balance tied into the available well control. Figure 8a and 8b show the prestack inversion for the P-impedance and Shear-impedance respectively for a profile through a 3D PSTM volume. This volume was created prior to applying any SOFD filtering on these data and used the best available gather conditioning, which included residual high-density velocity analysis and gather alignment. Indeed, the inversion produces a result that reasonably ties the well impedances as expected when the well data are filtered back with a high cut of 120-160 Hz. The inversion of the filtered data shown in Figures 8c and 8d, however, breaks out the impedance contrasts observed in the wells, particularly at the main horizons. In addition, the process has maintained the low-frequency component of the geology as can be seen in the impedance volumes. However, as should be expected, the filtered data does not generally tie the wells particularly well where the input data did not tie the wells.

Discussion

While the structure-oriented, frequency-dependent filtering process shows significant promise and the ability to substantially increase the useable bandwidth of seismic data, it is not without limitations. If we cannot obtain a stack section of sufficient quality to reliably compute structure tensors, then the noise attenuation will be limited. Even if the stack used to compute the tensors has a good signal-to-noise ratio, given enough noise in the data in a particular frequency band, the method will not be able to recover the signal. In addition, as the signal-to-noise increases, the reliability of the extracted signal amplitudes decreases and there comes a point at which the structure of the high-frequency component may be reliable, but the amplitudes may not be. Of course, this analysis of amplitude and frequency reliability can be done after the filtering process and the method and extent of balancing the spectrum can be chosen appropriately for the data and the intended attribute analyses to be performed.

While we did show that the low-frequency component can be enhanced, we did not show how this method would perform on very low frequencies, say 2-4 Hz. Indeed, although the bandpass filter used in Figure 2b was 2-3 Hz at the low end, there is very little signal in those data below 5 Hz and it would be extremely beneficial to be able to recover frequencies lower than this. So far we have only applied this method to seismic data acquired on land and with vibroseis as a source, so certainly some further testing on both land data acquired with dynamite and marine data would be prudent. However, our observations to date suggest that unlike the high frequencies, which appear to gradually reduce in signal-to-noise, the low frequencies in land seismic data can be enhanced up to a certain point and then we simply see very little evidence for any signal at all. This could be due to a number of factors in acquisition or processing or could be caused by the low-frequency energy being trapped in the near surface or primarily converted to surface waves.

In the examples shown in this paper, we were able to apply the filtering in a way that preserved the faults and subtle features observed in the stacked sections. However, increasing the

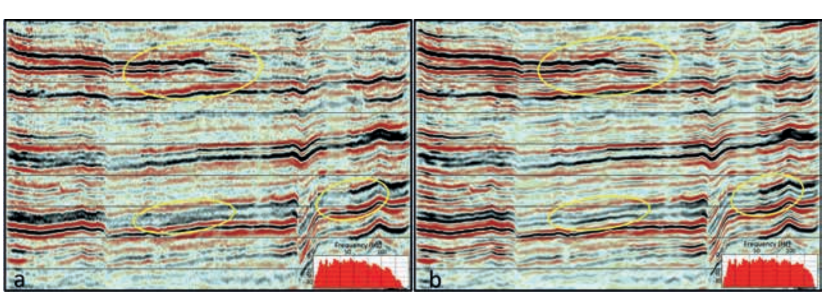

Figure 7 *Profiles through 3D stack volumes after applying pre-stack amplitude-preserving spectral balance. (a) Without SOFD filtering; (b) with SOFE filtering. As with the gathers shown in Figure 6, all filtering was applied in the surface X-Y domain with no filtering applied across offsets.*

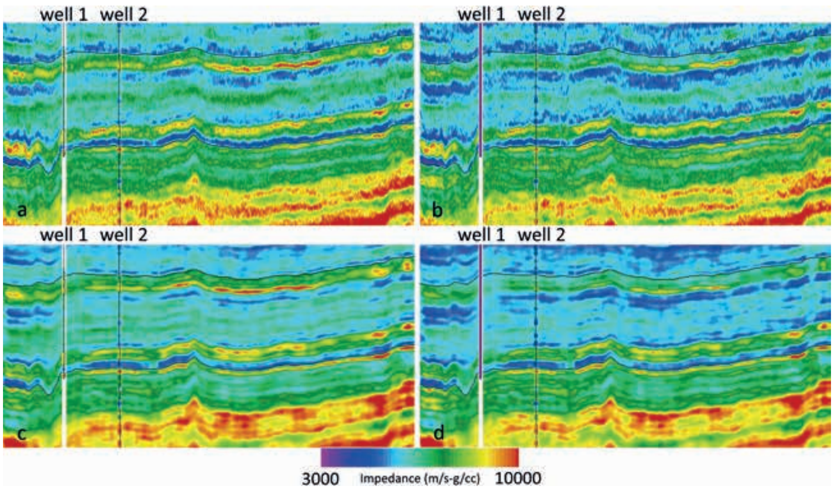

***Figure 8** Simultaneous inversion results for p-impedance (left) and s-impedance (right) with and without SOFD filtering. a) and b) P and S impedance without SOFD filtering; c) and d) P and S impedance with SOFD filtering.*

amount of filtering applied will reduce the lateral resolution to some extent. How this manifests itself is data-dependent. In some cases, for instance, the stack can strongly define the data discontinuities and applying the filtering increases the signal-to-noise sufficiently so that subtle features are easier to discern. However, for data that are generally noisy across a wide-frequency range, this may not be the case and a compromise between reducing noise and maintaining lateral resolution must be made as a function of frequency.

Conclusions

Using structure tensors to guide filtering of seismic data in a frequency-dependent manner can be used to effectively increase the useable bandwidth. The method does not make a priori assumptions about the geology/reflectivity or the relationship between the low and high-frequency components except that the frequency components have approximately the same physical structure. The method cannot force signal or noise to line up with the structure tensors, nor does it require that any signal be perfectly aligned with the structure tensors. Thus, features not observed in the data used to compute the structure tensors can be enhanced with the application of filtering using the structure tensors as a guide.

The output of the filtering process are data with improved signal-to-noise, particularly at the low and high frequencies. These data may then be used as input to processes which manipulate the relative amplitudes of different frequencies in the data without introducing intolerable noise and artifacts. Since the relative amplitudes of the frequency bands are maintained, the amount and type of frequency manipulation can be chosen based on the data characteristics and the final attributes to be computed independently of the filtering process.

Acknowledgements

The authors would like to thank ION Geophysical for supplying the seismic data used throughout this article and for supporting the work that led to the development of the SOFD technique described in this article and EnerVest for providing the well logs used in the SOFD analysis.

We would also like to thank NEOS for supporting continued development of the SOFD technique and for permission to publish.

Thanks also to Paul Constance for many helpful discussions.

References

Chopra, S. and Marfurt, K. [2008] Gleaning meaningful information from seismic attributes. *First Break*, **26**, 43-53.

Fehmers, G.C. and Höcker, C.F.W. [2003] Fast structural interpretations with structure-oriented filtering. *Geophysics*, **68**, 1286–1293.

Hale, D. [2009a] Image-guided blended neighbor interpolation of scattered data. *79th SEG Annual International Meeting, Expanded Abstracts*, 1127-1131.

Hale, D., [2009b] Structure-oriented smoothing and semblance, *Center For Wave Phenomena Consortium Meeting*. Colorado School of Mines, CWP-635.

Helmore, S., Merry, A.P. and Humberstone, I. [2007] Application of Frequency Split Structurally Oriented Filtering to Seismic Whitening and Seismic Inversion Workflows. *69th EAGE Conference & Exhibition, Extended Abstracts*, B027.

Hoeber, H.C., Brandwood S. and Whitcombe, D.N. [2006] Structurally consistent filtering. *68th EAGE Conference & Exhibition, Extended Abstracts,* G036.

Jenner, E. [2011] Combining VTI and HTI anisotropy in prestack time migration; workflow and data examples. *The Leading Edge*, **30**, 732-739.

Smith, M., Perry, G., Stein, J., Bertrand, A. and Yu, G. [2008] Extending seismic bandwidth using the continuous wavelet transform. *First Break*, **26**, 97-102.

Whitcome, D. and Hodgson, L. [2007] Stabilizing the low frequencies. *The Leading Edge*, **26**, 66-72.

Zhang, R. and Castangna, J. [2011] Seismic sparse-layer reflectivity inversion using basis pursuit decomposition. *Geophysics*, **76** (6), R147-R158.