# Gabor Deconvolution: Extending Wiener's method to nonstationarity

Gary Margrave, Linping Dong, Peter Gibson, Jeff Grossman, Dave Henley, and Michael Lamoureux

**SUMMARY**

Seismic data is arguably always nonstationary in its Fourier spectral character since anelastic attenuation processes are present everywhere. The major mathematical tool for dealing with this effect in seismic data processing is the Wiener deconvolution algorithm, which is at least a half century old. Paradoxically, this algorithm assumes that the seismic time series is stationary. Nevertheless, Wiener's algorithm deals with a limited class of attenuation effects in an average sense. When posed in the Fourier domain, the algorithm suggests its own generalization using a nonstationary extension of the Fourier transform such as the Gabor transform. A simple computational scheme for the Gabor transform is presented that is based on a set of windows that form a partition of unity. The use of the Gabor transform in a nonstationary deconvolution scheme is enabled because it factorizes the nonstationary convolutional model for a seismic trace with attenuation. The Gabor transform of a seismic signal is shown to be given approximately by the product of the Fourier transform of the source signature, the time-frequency attenuation function, and the Gabor transform of the reflectivity. Gabor deconvolution smooths the Gabor magnitude spectrum of the seismic signal to estimate the product of the magnitudes of the attenuation function and the source signature. A phase function is then estimated using the minimum phase assumption. The result is a complex-valued, time-frequency function that we call the spectrum of the propagating wavelet. When the original Gabor spectrum is divided by this time-frequency spectrum of the propagating wavelet, the result is an estimate of the Gabor spectrum of the reflectivity. Testing on nonstationary synthetic signals shows that Gabor deconvolution is dramatically effective in producing a very broadband reflectivity estimate, far better than Wiener's algorithm. Testing on real data also shows very highly resolved images from Gabor deconvolution. When compared with a "standard" image created using surface-consistent Wiener deconvolution together with time-variant spectral whitening, the Gabor images show subtle enhancement of detail and greater robustness in the presence of coherent noise and coverage gaps.

## INTRODUCTION

Norbert Wiener's deconvolution algorithm, brought to geophysics by the legendary Enders Robinson and Sven Treitel (Robinson and Trietel, 1967, and Peacock and Trietel, 1969) is both elegant and hugely important. In a simple prescription, it tells us how to whiten and phase-correct seismic data. Underlying this algorithm is the convolutional model of a seismic trace, which is usually stated by postulating that the seismic trace is the convolution of two other signals, one which we call the "wavelet" and another that we call the "reflectivity". Wiener deconvolution adopts this convolutional model and proceeds to make assumptions about both wavelet and reflectivity in order to estimate and remove (i.e. de-convolve) the former, thus revealing the latter. Plausible though this model is, it does not accommodate many physical effects such as attenuation and the general multiple series. For this reason, we have developed a nonstationary extension of the convolutional model and use it as the basis for a new deconvolution technique that directly estimates and removes both Q attenuation and source signature. The initial work on this new algorithm was presented at the 2002 CSEG National Convention (Margrave and Lamoureux, 2002).

Fundamental physical laws are usually formulated as partial differential equations (PDE's) and a particularly relevant one here would be a visco-elastic (or visco-acoustic) wave equation. We say "a" rather than "the" wave equation because there are many physical possibilities depending upon the details of our favourite physical model. A basic result from PDE theory, called Green's theorem, says that if we know the solution to any linear PDE for a perfectly impulsive source (i.e. a source with finite power concentrated at a single position and a single instant of time), then we can calculate the solution for a spatially extended source (like an array) or a source with temporal duration (like a dynamite blast) by convolution. The theory refers to the solution of the PDE for the impulsive source as the *impulse response* (or sometimes as the *Green's function*) and the description of the spatially and temporally extended source as the *source function*. To calculate the PDE solution for the extended source, we must convolve (in both space and time) the source waveform with the impulse response.

Green's theorem serves as a convenient point of departure to discuss the convolutional model as we use it in exploration seismology. First note that almost all of the interesting physics is contained in the impulse response. If we excite a perfectly impulsive source the resulting seismogram will contain all possible multiples, spherical spreading, attenuation effects, and virtually any other physical effect. The convolution of the impulse response with the source waveform is simply a spatial and temporal "blurring" of the impulse response.

Convolution is a familiar operation to geophysicists and we will not review it in detail here; however, it is relevant to discuss the *stationarity* property of convolution. For simplicity, assume the source waveform is extended in time only with a particular waveform and remains impulsive in space. Then the convolution is in time only and consists of the process of replacing each sample of the impulse response with the extended source waveform scaled by the seismogram sample. Thus the seismogram time series is visualized as a "scaled superposition" of source waveforms. The convolution

process is said to be stationary because the source waveform is the same for all samples of the impulse response regardless of their arrival time.

Convolution is an example of a *replacement process*, which means that each sample in a time series is being replaced by another time series (scaled by the sample). As a simple illustration, consider the time series s=[1, 0, -.5, .2, 0, 0] and let us replace each sample by $s_k$[1, -.1]=[$s_k$, -.1$s_k$], where $s_k$ means the $k^{th}$ sample. After replacing only the first sample we have s=[1, -.1, -.5, .2, 0, 0] and, in fact, this is also the result after replacing the second sample since it is zero. Replacing the third sample is slightly different since it requires that we sum the -.1$s_k$ term into the fourth sample. This gives s=[1, -.1, -.5, .25, 0, 0]. Now, we replace the fourth sample (being sure to use the original value of sample four which was .2 not .25) to get s=[1, -.1, -.5, .25, -.02, 0]. This particular illustration of a replacement process since the *replacement function*, [1, -.1], was the same for each sample. We can easily imaging a *nonstationary* replacement process where each sample is replaced by a different function. For instance, we could insert a number of zeros between 1 and -.1 that was dependent upon the position of the sample we are replacing. So, we might replace the first sample with $s_1$[1, -.1], the second with $s_2$[1, 0, -.1], the third with $s_3$[1, 0, 0, -.1] and so on. We can such a nonstationary replacement process a *nonstationary convolution*.

If Green's theorem were to be the basis for deconvolution theory, then all deconvolution would achieve would be the estimation of the impulse response. We would then be left with a signal that still contains all possible multiples, attenuation effects, and other physics. This is clearly not sufficient progress as the colloquial expectation is that the result from deconvolution is an estimation of the earth's *reflectivity*, a term usually taken to mean a time series whose samples are the primaries-only, normal-incidence, reflection coefficients positioned at their two-way traveltimes. Is it somehow possible to contrive an alternative convolutional model in which the multiples and attenuation are part of the source waveform rather than the impulse response? Put another way, since it is reflectivity that we want to estimate, can we decompose the impulse response into a reflectivity that is altered by operations of attenuation and multiple contamination and then move the multiples and attenuation into the source waveform term to create an "effective" or "equivalent" source waveform that is convolved directly with reflectivity?

In the general case, the answer to this question is a definitive no. The limiting consideration is that the convolution process is necessarily stationary while multiples and attenuation are clearly nonstationary. To appreciate this fundamental nonstationarity, consider the situation of multiples first and for simplicity let attenuation be absent. As the leading seismic pulse propagates down into the earth, it is accumulating an ever-growing tail of multiples. Every interface that the wave encounters gives rise to both reflections and transmissions of both P and S waves. The upward traveling reflection of the leading wave will eventually be recorded as a primary at the geophone. However, as it moves up it generates downward traveling reflections off each interface that follow the leading pulse down. Thus we can view any particular interface as being impacted by the leading pulse followed by a great many secondary pulses of many different sizes and polarities. The longer the leading pulse propagates, the longer is this *multiple train*.

Since the reflectivity should contain only the leading pulses (primaries) the calculation of the impulse response from the reflectivity requires the replacement of each sample of the reflectivity with a progressively growing multiple train. This is not a stationary process since the replacement function is growing continually with time.

The consideration of attenuation is similar. All earth materials show some degree of attenuation and the attenuation increases with the distance that the wave travels. Each sample of the impulse response has a unique traveltime and that may correspond to a great many travelpaths in the subsurface. Now with attenuation included, the calculation of the impulse response from the reflectivity requires replacing each reflectivity sample with a multiple train that becomes progressively more and more attenuated. Thus attenuation is also a nonstationary process.

So, this analysis begs the question, how can Wiener deconvolution possibly work? Have we been deceiving ourselves for the past forty years? Thankfully, it is easy to see that we have not and that Wiener's algorithm does work in a useful, though approximate, sense. If we drop the notion of including all possible multiples as part of an effective source waveform, then perhaps there are selected multiples that can be considered as a stationary modification of the reflectivity. An important case that is essentially stationary is the case of the first-order water layer multiple at zero offset. Neglecting all other multiples, the water layer multiple is essentially a copy of the source waveform that lags behind the leading pulse by the two-way traveltime in the water layer. Thus we can calculate an artificially abbreviated impulse response that includes only the effects of the first-order water layer multiple, as a convolution of the reflectivity with an equivalent source that has two impulses. The first impulse is of unit amplitude at time zero and represents the leading pulse while the second is at the two-way water-layer traveltime and has an amplitude equal to the negative of the water bottom reflection coefficient.

The key feature that allows us to model the water layer multiple as stationary is that the water layer is above all other reflectors. A pulse that reflects off a deeper interface must traverse the water layer twice, on the way down and up. Thus each deeper reflector is impacted with an effective waveform that has the same modified form for the water layer. This suggests that any interface that generates significant multiples and is above the reflectors of "interest" can be similarly treated. Since it is common practice in Wiener deconvolution to select a "design window", perhaps 1 second or so long, that encompasses the target horizons, we anticipate that multiples from horizons above the design window can be modelled as stationary and included in the effective source. This process cannot be extended to all possible multiples and retain the stationarity property because the $n^{th}$ reflector is impacted by an effective wavelet containing multiples from the upper n-1 reflectors. Thus, in general, each reflector experiences a different multiple train.

For attenuation, most of the effect commonly happens in the unconsolidated near-surface and so a good first order approximation is that each deeper reflector is impacted by the same attenuated source waveform. This approximate model allows us to consider some of the earth attenuation effects as stationary and so lump them into an effective source. However, as soon as we admit that the real situation is that attenuation occurs all along the raypath, then we are forced to concede that each reflector experiences a

uniquely attenuated incident waveform. Thus, attenuation is also nonstationary in the general case.

So we must conclude that, if multiples and attenuation that originate from all possible positions along the travelpaths of waves are to be included in the seismic trace model, then that model must be a nonstationary one. In the next section we introduce such a nonstationary model for the inclusion of attenuation. The case of multiples is left for the future.

## A NONSTATIONARY CONVOLUTION MODEL

Figure 1 illustrates the stationary convolutional model by showing a source waveform being convolved with a reflectivity time series. Thus we are modelling no multiples or attenuation other than those that can be considered as part of the effective source. The convolution is achieved by constructing a special matrix containing the source waveform and multiplying it into a column vector containing the reflectivity. In the convolution matrix, the source waveform appears in each column with its first sample shifted to lie along the main diagonal. This results in a matrix where every diagonal (from upper left to lower right) has a constant entry. This symmetry is called *Toeplitz* and is the telltale sign of a stationary process.

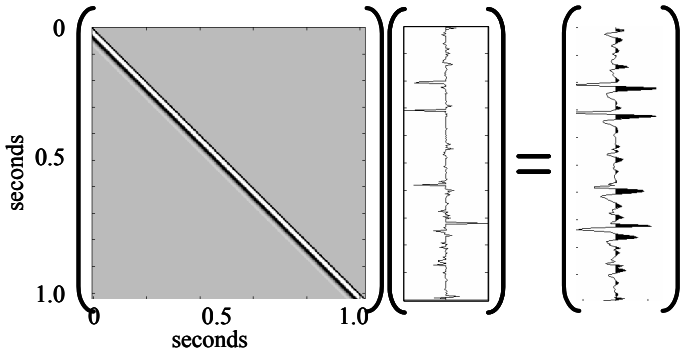


FIG. 1. The stationary convolutional model for a seismic signal is depicted as the multiplication of a column vector containing the reflectivity by a Toeplitz matrix containing the source waveform

In contrast, Figure 2 shows a nonstationary replacement process in which the same reflectivity as Figure 1 is being acted upon by a matrix that represents a nonstationary attenuation process. Each column of the nonstationary convolution matrix contains the effective source for the time corresponding to the main diagonal. From upper left, representing early times, to lower right, representing later times, the effective source waveform shows a steady decay in both amplitude and frequency due to the attenuation mechanism. When applied to the reflectivity column vector, the result is a very different seismogram from Figure 1 and shows a progressively attenuated response. This is a nonstationary convolution model that is formulated to include the effects of attenuation wherever it may occur.

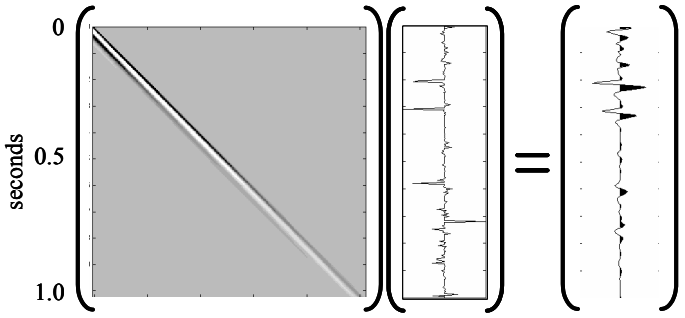


FIG. 2. The nonstationary convolutional model proposed in this paper is depicted as a matrix-vector multiplication similar to that of Figure 1. The nonstationarity is seen in the matrix where the source waveform steadily attenuates as it moves to later times.

The matrix operation shown in Figure 2 is most conveniently expressed as a mathematical for known as a *pseudodifferential operator*. This somewhat intimidating term is really nothing more than a nonstationary filter that proceeds exactly as discussed in the preceding section. Formally, we can represent Figure 2 by the equation

$$\hat{s}(f) = \hat{w}(f) \int \alpha(t,f) r(t) e^{-2\pi i f t} dt \,. \tag{1}$$

In this equation, the integral is taken over all time, $t$ and $f$ are time and frequency, $r$ is reflectivity, $w$ is the source waveform, $s$ is the seismic trace, a "hat" over a variable indicates the (forward) Fourier transform, and $\alpha$ is the time and frequency dependent attenuation function. In the case of the commonly assumed "constant Q" model, $\alpha$ is given by

$$\alpha(t,f) = \exp\left(-\frac{\pi t}{Q(t)}\left(f + iH(f)\right)\right) \tag{2}$$

in which $H$ denotes the Hilbert transform. The stationary convolution model emerges as a special case of equation (1) corresponding to infinite $Q$. In this case $\alpha = 1$ and the integral in equation (1) becomes an ordinary Fourier transform giving

$$\hat{s}_\infty(f) = \hat{w}(f)\hat{r}(t), \tag{3}$$

which we can recognize as the standard convolution model since multiplication in the frequency domain is convolution in the time domain. (The subscript on $s$ in equation (3) is to distinguish the stationary signal from the nonstationary by reminding us that the former corresponds to $Q = \infty$.)

Equation (1) is structured to prescribe the Fourier transform of the seismic trace rather than the trace in the time domain because it is mathematically simpler. When this equation is rewritten in the time domain and specialized to discretely sampled time series with particular values for the parameters, the representation shown in Figure (2) results.

## WIENER'S ALGORITHM

We now summarize Wiener's algorithm for solving the stationary deconvolution problem posed by the stationary trace model of equation (3). We choose to present this model in the frequency (i.e. Fourier) domain because our generalization to the nonstationary Gabor world starts from there. Robinson and Trietel usually presented it in the time domain but it is acompletely equivalent to our frequency domain algorithm. (We will indicate the equivalent time-domain steps by parenthetical notes like this.)

A key ingredient of the method are assumptions made about $w(t)$, the source waveform, and $r(t)$, the reflectivity function. The former is assumed to be *minimum phase* while the latter is taken to be *white*. These phrases have specific mathematical meanings. For example, minimum phase can be shown to mean that the wavelet has the least possible phase delay (group) of all causal, invertible wavelets with the same amplitude spectrum. (Causal wavelets are those which vanish for negative times $t < 0$.) We mention this definition first because is suggests the meaning of the term *minimum phase*. There are other equivalent definitions including (1) a causal, stable wavelet with a causal stable inverse is minimum phase (stable simply means finite energy), (2) a wavelet whose Fourier phase spectrum is the Hilbert transform of the logarithm of its Fourier amplitude spectrum is minimum phase. These definitions are equivalent in that any one implies the others. Of particular importance in these statements is the implication that a minimum-phase wavelet can have no zeros in its amplitude spectrum, since if it did we could not invert it nor could we take the logarithm of the amplitude spectrum. Since no physical system can possibly radiate at all frequencies from zero to infinity, it is generally true that physical systems can only be approximately minimum phase. Another complexity is that the digital recording of a physical minimum-phase system may not be minimum phase due largely to the effects of the anti-alias filter. We ignore these problems.

The assumption of *white* reflectivity means that the reflectivity is assumed to be a stochastically random time series. Such time-series can be shown to have two equivalent properties: (1) their autocorrelation is everywhere zero except for a single impulse at zero lag and (2) their power spectrum is constant at all frequencies. The latter property implies the name since white light has equal power at all visible frequencies. Even with synthetic reflectivities, this can only be realized approximately since it can be shown that a random signal with a truly white power spectrum must be infinitely long. Furthermore, with real well logs, this property is never precisely true. The autocorrelation property means that a white reflectivity must be statistically uncorrelated with itself. This is never true for geologic systems which might correlate strongly over small time intervals even if they are nearly uncorrelated over larger intervals. If a geologic system that deposits shale exists at a particular moment of time, then at a neighboring moment, the probablility is high that shale will also be deposited. At moments more removed in time, nearly anything else may happen; hence, geologic correlations tend to be short. A time series with strong correlations at small lags and randomness at large lags tends to have less power at low frequencies than at high frequencies and is said to be *blue*.

Nevertheless, assuming these properties of the source waveform and the reflectivity allows a solution to the deconvolution problem posed by equation (3). The first step is to take the magnitude of the complex-valued Fourier spectra, thus discarding all phase information. The resulting amplitude spectra for a typical case with synthetic data is shown in Figure 3. Wiener's insight was to suggest that the general shape of the amplitude spectrum of the seismic signal comes entirely from the spectral shape of the source waveform. Put another way, if we smooth the amplitude spectrum of the seismic signal, we will estimate the amplitude spectrum of the source waveform. This step requires the assumption of white reflectivity. For a blue reflectivity, some of the spectral shape of the seismic amplitude spectrum comes from the reflectivity. (In the time domain, Wiener actually said autocorrelation of the seismogram is similar to the autocorrelation of the source waveform. If we window the autocorrelation, this is analogous to smoothing the power spectrum, then we estimate the autocorrelation of the source waveform.) There are infinitely many ways to smooth the amplitude spectrum and hence there are a great many practical questions in a successful Wiener algorithm.

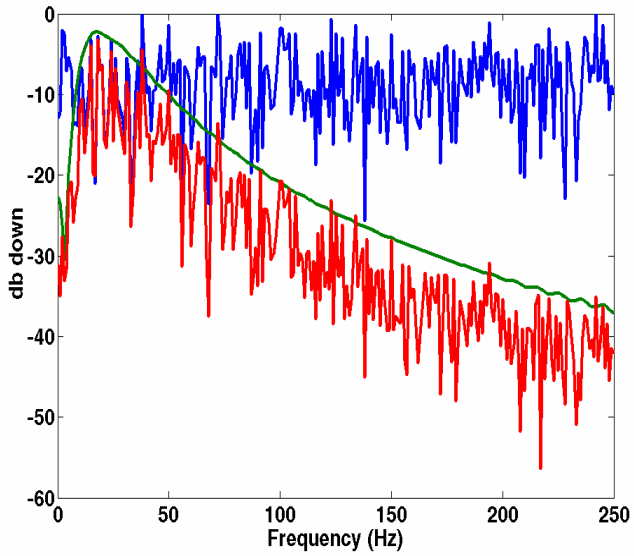


FIG. 3. The stationary convolutional model is depicted in the Fourier domain. Blue is the Fourier amplitude spectrum of the reflectivity, green is that of the wavelet, and red is that of the seismic signal. Wiener deconvolution assumes that the red curve is measured and can be smoothed to estimate the green curve. The spectral division of red by green gives blue.

One we have estimated the amplitude spectrum of the source waveform, we need to estimate the waveform's phase spectrum. This is where the assumption of minimum phase comes in. In the frequency domain, we calculate the phase spectrum as the Hilbert transform of the logarithm of the amplitude spectrum. This is problematic for two reasons (1) the amplitude spectral estimate may contain zeros that will prevent us from taking the logarithm, and (2) the Hilbert transform requires an integration over all frequencies to estimate the phase spectrum at any frequency and the amplitude spectral estimate will surely be wrong (noise dominated) at high frequencies. Both of these problems are dealt with in practice by adding a small constant to the amplitude spectral estimate at all frequencies. Called a stability constant or white noise term, this is usually chosen to be larger than the amplitude spectral estimate at any frequencies judged to be

random noise. (In the time domain, the minimum-phase assumption is invoked during the solution of the *normal equations*. This refers to a system of equations whose solution gives the inverse filter for the source waveform knowing only its autocorrelation. The white noise term is a small modification to the zero-lag of the autocorrelation that ensures that the normal equations have a solution.)

Having estimated both the amplitude and phase spectra of the source waveform, we can now remove the source waveform from the seismic signal. This is the actual mathematical process of deconvolution and we accomplish it by simple division in the Fourier domain. That is, we simply divide the Fourier spectrum of the seismic signal by the Fourier spectrum of the estimated source signature. This spectral division should always be numerically stable since we have already guaranteed that the amplitude spectrum of the source waveform estimate has no zeros. (In the time domain, the inverse of the source signature is directly estimated and deconvolution is accomplished by convolving this inverse with the seismogram.)

The result from the spectral division is an estimate of the Fourier spectrum of the reflectivity. Especially when accomplished by spectral division, this estimate tends to be very broadband and will have nearly constant power at all frequencies. In general, this will be too much "whitening" and a deconvolution is often followed by a bandpass filter that restricts the whitening to some assumed signal band. Perhaps a better alternative is to allow the reflectivity estimate to remain broadband until stack and use the noise attenuation property of stacking to attenuate incoherent noise.

A final feature of the practical implementation of the Wiener algorithm is the use of a *design gate*. This means that a particular time zone of the seismic trace is designated as the zone for which the source waveform will be estimated. This estimated waveform is then used to deconvolve the entire trace. A direct implication is that, while the estimated source waveform may be optimal in the design gate, it may be very inappropriate elsewhere. A common result is that a source waveform estimate from a middle time zone leads to over-whitening for shallower times and under-whitening at deeper times. This is sometimes dealt with but using multiple design gates for a single trace.

A related problem is that it can be difficult to find a design gate on certain traces that is not contaminated by unwanted coherent source noise. A source waveform from such a time zone can be inappropriate everywhere, even in the design gate. This has lead to a significant evolution of the Wiener algorithm called *surface-consistent deconvolution*. (The algorithm just described is now called *single-channel deconvolution*.) In this technique, each trace has at least two waveforms estimated for it, one that is said to be source consistent and another that is receiver consistent. In a simple version of the technique, the amplitude spectra (from the design gate) the ensemble of all traces common to a given source are averaged and a single source waveform is estimated. This process is repeated for the ensemble all traces common to a given receiver to estimate a receiver waveform. In more sophisticated approaches, the source and receiver waveforms are designed via least squares rather than simple ensemble averaging. Today, surface-consistent deconvolution has largely replaced single-channel Wiener deconvolution largely because the former produces fewer artefacts related to source-generated noise.

## THE GABOR TRANSFORM

We now turn to the solution of the nonstationary deconvolution problem as posed by equation (1). An essential tool in our solution algorithm is an extension of the Fourier transform known as the Gabor transform. Originally proposed by Dennis Gabor, the inventor of the hologram, as a tool in quantum mechanics, the Gabor transform uses a window function to localize a signal to a certain time range and then applies the Fourier transform. Sometimes called a *short-time Fourier transform*, the Gabor transform uses a suite of windows designed to allow the process to be invertible. Gabor initially suggested using Gaussians as window functions but the theory has evolved since then to allow a great variety of windows.

We have developed a particular implementation of the Gabor transform, appropriate for discretely sampled, bandlimited signals based on the concept of a partition of unity (POU). A POU is a set of windows defined on the real line and so chosen that they always sum to unity (Figure 4). If $\Omega(t) \geq 0$ denotes a non-negative window function, then it forms a POU if

$$\sum_{n\in\mathbb{Z}} \Omega(t - n\Delta\tau) = \sum_{n\in\mathbb{Z}} \Omega_n = 1 \tag{4}$$

where by $\Omega_n$ we mean the window $\Omega$ translated (shifted) to the time $n\Delta\tau$. The POU defined by equation (4) uses a single window function translated at regular intervals along the real line to form a POU. This is simple and convenient but not necessary. The windows of a POU need not be formed from translations of a "mother" window, they just need to sum to unity.

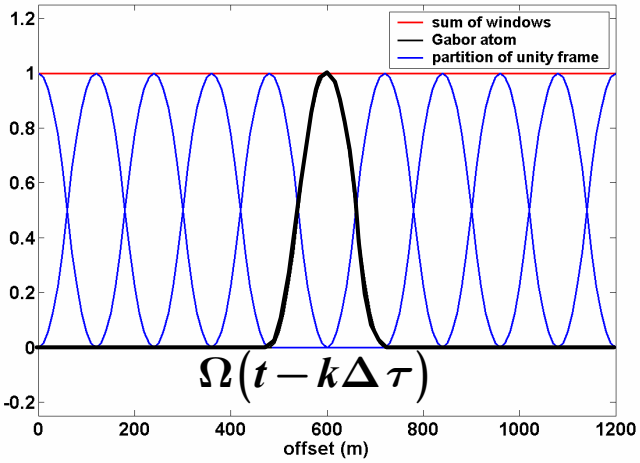


FIG. 4. A set of windows that form a POU (partition of unity) is shown.

Next we define an *analysis window*, $g$, and a synthesis window, $\gamma$, by the equations

$$g_n = \Omega_n^p, p \in [0,1], \tag{5}$$

and

$$\gamma_n = \Omega_n^{1-p}. \tag{6}$$

So $g$ and $\gamma$ are just fractional powers of $\Omega$ such that

$$\sum_{n\in\mathbb{Z}} \gamma_n g_n = \sum_{n\in\mathbb{Z}} \Omega_n = 1. \tag{7}$$

We define the *Gabor slice* using the analysis window $g$ as

$$s_n(t) = s(t)\, g_n(t) = s(t)\, g(t - n\Delta\tau) \tag{8}$$

and the Gabor transform is the Fourier transform over the set of Gabor slices

$$G[s](n,f) = \widehat{s_n}(f). \tag{9}$$

This equation emphasizes that the Gabor transform of $s$ is a two-dimensional time-frequency decomposition. Here time is denoted by the integer $n$ and frequency by the real number $f$. The use of a continuous variable for frequency is just a notational convenience as, in any numerical implementation we use the discrete Fourier transform and calculate equation (9) over a discrete set of frequencies as well as times.

The inverse the Gabor transform is easily accomplished using the synthesis window $\gamma$, the inverse Fourier transform, $F^{-1}$, and summation over $n$, by the equation

$$G^{-1}\left[G(s)\right] = \sum_{n\in\mathbb{Z}} \gamma_n F^{-1}\left[\widehat{s_n}\right] = s. \tag{10}$$

That this recovers $s$ exactly follows because the inverse Fourier transform recovers the Gabor slice (equation (8)) and the summation over $n$ using equation (7) recovers $s$.

While $G^{-1}G = I$ ($I$ is the identity operator), $GG^{-1} = P$ is not the identity but is rather a projection operator, $P$, from functions defined on the entire 2D time-frequency plane, $\mathbb{R}\times\mathbb{R}$, onto the range of the forward Gabor transform. This contrasts with the case of the Fourier transform which yields the identity for both cases. The Fourier transform uses the orthonormal basis of $\cos(2\pi ft) + i\sin(2\pi ft)$ to represent a signal. $F^{-1}F = I$ means that the signal can be recovered identically from its Fourier transform while $FF^{-1} = I$ means that each signal has a unique Fourier transform. The Gabor transform is not a decomposition into an orthonormal basis, instead the windowed sines and consines form a more general construct called a *frame*. A frame can be used like an orthonormal basis to decompose a function (i.e. we can use a frame to calculate a spectrum for any signal) but the decomposition (i.e. spectrum) is not generally unique.

Another way of viewing this situation is to let $L^2(\mathbb{R})$ and $L^2(\mathbb{R}\times\mathbb{R})$ denote the spaces of 1-dimensional and 2-dimensional, complex-valued, square integrable functions. Then our seismic signals live in $L^2(\mathbb{R})$ and $F$ maps each signal to another point in $L^2(\mathbb{R})$. In fact any point in $L^2(\mathbb{R})$ can be regarded as either a signal or a Fourier

spectrum. In contrast, only a subset of $L^2(\mathbb{R}\times\mathbb{R})$ can be regarded as the Gabor transform of some signal in $L^2(\mathbb{R})$. This is because the well-known *uncertainty principle* imposes smoothness requirements upon the Gabor spectrum. Thus, for any choice of window functions, there will always be functions in $L^2(\mathbb{R}\times\mathbb{R})$ that cannot be the Gabor transform of any function in $L^2(\mathbb{R})$. These will be functions of time and frequency that change more abruptly than is allowed by the uncertainly principle. If we insist on computing the inverse Gabor transform of such a function, we will indeed calculate a signal in $L^2(\mathbb{R})$ but when we then calculate the forward Gabor transform of this signal we will not be back where we started. Instead we will have *projected* the original function with its too-rapid variation onto the range of the forward Gabor transform. That is, we will have calculated a smoothed version of the original function.

An important notion in the theory of the discrete Gabor transform is that of the *redundancy* of the representation of equation (9). For an $N$ point time series, the redundancy of the Gabor transform is defined as the ratio of the number of samples required in the time-frequency representation to N. The discrete Fourier transform can be viewed as an end-member case with a redundancy on 1. For the case shown in Figure 4 where the windows of the POU are compactly supported (meaning they vanish outside a compact interval) and the overlap between adjacent windows is 50%, the redundancy can be shown to be 2. By increasing the window length while decreasing the overlap, the redundancy can be made to approach, but never equal, 1. (It might seem that we could achieve a redundancy of 1 with "boxcar" windows of zero overlap; however, some overlap is desirable for filtering and this special case is of little practical interest.)

Figure 5 illustrates the Gabor spectrum of a synthetic seismogram constructed using the nonstationary trace model of equation (1). The Gabor spectral plot shows only the magnitude of $G[s]$ and black indicates a large positive value while neutral grey is a very small number. The Gabor spectrum shows clear decay in both time and frequency. Also evident are the imprint of the reflectivity and the bandlimiting due to the source waveform.

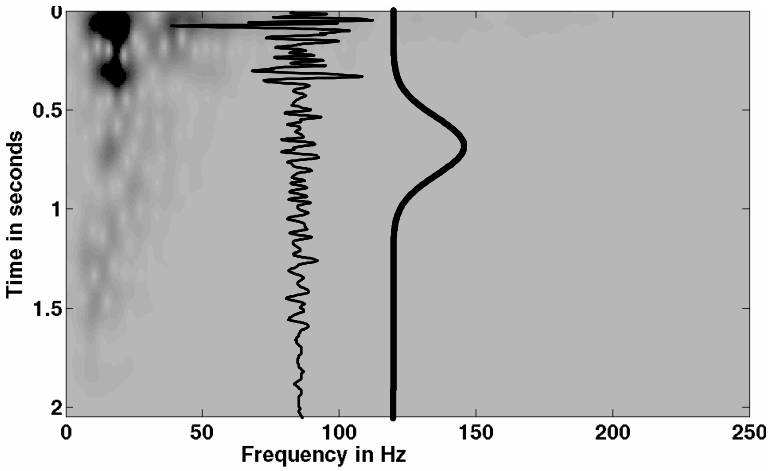


FIG. 5. A nonstationary synthetic seismic signal and a typical Gabor window are shown superimposed on the Gabor magnitude spectrum of the seismic signal. The Gabor plot has time on the vertical axis and frequency on the horizontal axis. High amplitudes are darker greys.

**Factorization of the Nonstationary Trace Model**

In a complicated technical result that we will not reproduce here, we have shown that when the Gabor transform is applied to the nonstationary trace model proposed in equation (1), the result is approximately a product of factors that is analogous to equation (3), given by

$$G[s](n,f) = \hat{w}(f)\alpha(t = n\Delta\tau, f)G[r](n,f) + \cdots \tag{11}$$

where the ellipses indicates the presence of further terms required to make equation (11) an equality. (Since equation (1) gives the Fourier transform of our trace model, we must first apply the inverse Fourier transform to equation (1) before applying the forward Gabor transform.)

Though a formal technical justification of equation (11) is inappropriate here, we can see why it is reasonable by a physical-mathematical argument. The Gabor slice of equation (1) is given by

$$s_n = g_n s = g_n F^{-1}\left[\hat{w}(f)\int \alpha(t,f)r(t)e^{-2\pi ift}dt\right] \tag{12}$$

where the general time dependence has been dropped from the notation. Now, if the analysis window, $g$, is sufficiently local that the attenuation function changes little across the window, we can replace $\alpha(t,f)$ by $\alpha(n\Delta\tau, f)$, that is we replace the attenuation function with its value at the window centre. Then, $\alpha(n\Delta\tau, f)$ can be pulled out of the integral in equation (12), which then becomes an ordinary Fourier transform and we have

$$s_n \approx g_n F^{-1}[\hat{w}_n \hat{r}] \tag{13}$$

where $\hat{w}_n(f) = \hat{w}(f)\alpha(n\Delta\tau, f)$ is the effective waveform at the window centre and we have dropped the frequency dependence form the notation as well. Now, since the Gabor transform is just the Fourier transform of the Gabor slice, we have

$$G[s] \approx F\left[g_n F^{-1}[\hat{w}_n \hat{r}]\right] = \hat{g}_n \bullet (\hat{w}_n \hat{r}) \tag{14}$$

where $\hat{g}_n \bullet$ means convolution with $\hat{g}_n$ and the last step follows because windowing in the time domain is convolution in the Fourier domain with the Fourier transform of the window. Equation (14) has the operations of multiplication by $\hat{w}_n$ followed by convolution with $\hat{g}_n$ being applied to the Fourier transform of *r*. If $\hat{g}_n$ is sufficiently narrow with respect to $\hat{w}_n$, then we can commute these operations. Formally we write

$$G[s] \approx \hat{g}_n \bullet (\hat{w}_n \hat{r}) = \hat{w}_n (\hat{g}_n \bullet \hat{r}) + [\hat{g}_n \bullet, \hat{w}_n] \hat{r} \tag{15}$$

where $[\hat{g}_n \bullet, \hat{w}_n]$ is called a commutator and is just the difference between applying the operations in one order and then the other. If we neglect the commutator in equation (15) then we have the result we want

$$G[s] \approx \hat{w}_n (\hat{g}_n \bullet \hat{r}) = \hat{w}_n F[gr] = \hat{w}_n G[r]. \tag{16}$$

Given the definition of $\hat{w}_n(f) = \hat{w}(f)\alpha(n\Delta\tau, f)$, this is equivalent to the statement in equation (11).

Though not fully rigorous this derivation makes clear where the approximations in the factorization of equation (11) lie. First, we had to assume that the variation of the attenuation function over $g_k$ could be approximated by its value at the window midpoint. Since low values of *Q* mean rapid attenuation, this links the *Q* value with the largest acceptable window length, the higher the *Q* value the larger this upper bound on window length can be. Second, we had to commute convolution with $\hat{g}_n$ multiplication by $\hat{w}_k$. There are two end-member cases where this commutation can be done without error. If $\hat{g}_n$ is a Dirac delta function (or a unit spike in the discrete case) then $\hat{g}_n \bullet$ commutes with any $\hat{w}_n$. This is only precisely true if $g_n = 1$ everywhere, but we surmise that there will be some lower bound on window length below which $\hat{g}_n$ becomes too broad and the error expressed by $[\hat{g}_n \bullet, \hat{w}_n]$ become unacceptable. Also, if $\hat{w}_n$ is a constant over all frequencies, then it commutes with any $\hat{g}_n \bullet$. While we might expect $\hat{w}_n$ to vary smoothly it is certain not to be constant. Intuitively, we might expect the commutator to be small provided that $\hat{w}_n$ is nearly constant over the width of $\hat{g}_n$ though we lack a proof at this time. Equivalently, we might also expect the commutator to be small if, in the time domain, $g_n$ is much more broad than $w_n$. To summarize, the first approximation places an upper bound, determined by the temporal variation of $\alpha(t, f)$, on the allowable window size while the second approximation places a lower bound, determined by the temporal length of $w_n$, on the allowable window size.

While it is difficult to make the window length bounds precise, we have conducted a great many numerical experiments on synthetic data and find that the factorization approximation of equation (11) holds for a broad range of Q values and window sizes.

The factorization property is the key conceptual ingredient in our deconvolution algorithm that we now detail.

**The Gabor Deconvolution Algorithm**

We now assume that equation (11), without any implied error terms, is a precise description of the Gabor transform of a seismic signal. Furthermore, we will find it convenient to assume that the attenuation function is of the form given by equation (2). This is an assumption of minimum phase since equation (2) has the property that the phase of $\alpha(t,f)$ is given by the Hilbert transform of the logarithm of the magnitude, $|\alpha(t,f)|$, where the Hilbert transform is taken over *f* at constant *t*. We also assume that the source waveform, $w(t)$, is minimum phase.

A virtue of the nonstationary spectral factorization of equation (11) versus the stationary factorization of equation (3) is that the former has a clear separation of source waveform and the attenuation process. If equation (3) is to be valid then the source waveform must include an attenuation term, but there is no apparent way to separate these. In the nonstationary factorization, the source waveform is a function of *f* only while the attenuation function depends upon both *t* and *f*. In fact, if *Q* is constant, then according to equation (2), $|\alpha(t,f)|$ should be constant on the hyperbolic family $tf = \text{constant}$.

As in stationary deconvolution, we also assume that the reflectivity is a random sequence though precisely how this manifests in the Gabor domain is not clear. In a practical statement, we assume that the detail in $G[s](n,f)$ is due to $G[r](n,f)$ while the general trends are due to $\hat{w}(f)\alpha(n,f)$. This implies that there exists some smoothing or fitting operation that can separate $G[s](n,f)$ into detail and trend and so estimate $\hat{w}(f)\alpha(n,f)$ and $G[r](n,f)$.

As with Wiener deconvolution, the first step in our algorithm is to discard the phase by taking the magnitude of $G[s](n,f)$. We then amply some sort of smoothing operation whose goal is to eliminate all features attributable to $G[r](n,f)$. We have experimented with various smoothers and least squares fitting. The simplest possible smoothing procedure is to convolve the Gabor magnitude spectrum of the seismic signal with a 2D boxcar, whose time and frequency dimensions must be chosen empirically. A major problem with this approach is that it has an amplitude equalization effect, much like an AGC, that tends to distort the estimated reflectivity. Usually better is a technique we call *hyperbolic smoothing*. Since the level-lines of the attenuation function, for *Q* a constant, are the hyperbolic family defined by *tf* = constant, we are motivated to calculate the mean value of the Gabor magnitude spectrum along such contours. In this case, we take the mean along hyperbolic contours as an estimate only of $|\alpha(n,f)|$. We then divide $|G[s](n,f)|$ by the estimate for $|\alpha(n,f)|$. The result of this division is averaged over t and smoothed in f with a convolution operator to estimate $|\hat{w}(f)|$. Another possibility is to perform a least-squares fit of the model of equation (11) to $G[s](n,f)$.

In any case, assume for the moment that a particular smoothing or fitting operation has been conducted on $G[s](n,f)$. This becomes an estimate of the Gabor magnitude spectrum of the propagating wavelet that we will call $|\hat{w}_\alpha|$ and we assert that

$$|\hat{w}_\alpha|(n,f) \approx |\hat{w}(f)\alpha(n,f)| \tag{17}$$

where the approximation holds with a hopefully small, but admittedly unknown, error term. We then invoke the minimum-phase assumption to calculate the phase of $\hat{w}_\alpha$. With the inclusion of a small stability constant $\varepsilon$, as in Wiener deconvolution, we estimate the entire Gabor spectrum of the propagating wavelet as

$$\hat{w}_\alpha = \left(|\hat{w}_\alpha| + \varepsilon\right) e^{iH\left[\ln\left(|\hat{w}_\alpha|+\varepsilon\right)\right]} \tag{18}$$

where we understand that both sides of the equation depend on discrete time *n* and frequency *f*, and the Hilbert transform is to be taken over frequency.

Having completed an estimation of the Gabor spectrum of the propagating wavelet, we now perform the actual deconvolution by spectral division in the Gabor domain

$$G[s_d](n,f) = \frac{G[s](n,f)}{\hat{w}_\alpha(n,f)} \tag{19}$$

where we have used $s_d$ to symbolize the deconvolved seismic trace and the division is understood to be a point-by-point operation using complex arithmetic. $s_d$ itself is recovered by an inverse Gabor transform of the result from equation (19). Usually, it is convenient to bandlimit $G[s_d](n,f)$ in the Gabor domain prior to the inverse transform. Since the attenuation function may be nearly constant along the hyperbolae *tf* = constant, we find it suitable to define a high frequency cutoff whose value follows a particular hyperbolic trajectory.

We now begin an extended example of the process beginning with the synthetic signal shown in Figure 5. Figure 6 shows the estimate of $|\hat{w}_\alpha|$ calculated by boxcar smoothing of the magnitude of the spectrum of Figure 5. The boxcar had dimensions of .5sec x 20Hz. Figure 7 shows the corresponding estimate of $|\hat{w}_\alpha|$ calculated by hyperbolic smoothing. Though perhaps appearing similar, the estimate of Figure 7 is generally superior to that of Figure 6. A noticeable problem with the latter is occurs near 20 Hz between .8 and 1.3 seconds. There is a significant amplitude low near 1.1 seconds and amplitude then increases near 1.3 seconds. This is a physical impossibility since it predicts that the propagating wavelet somehow gained energy. It is actually caused by incomplete removal of reflectivity effects during the smoothing process. We do not illustrate the minimum phase spectrum computed from these magnitude spectra.

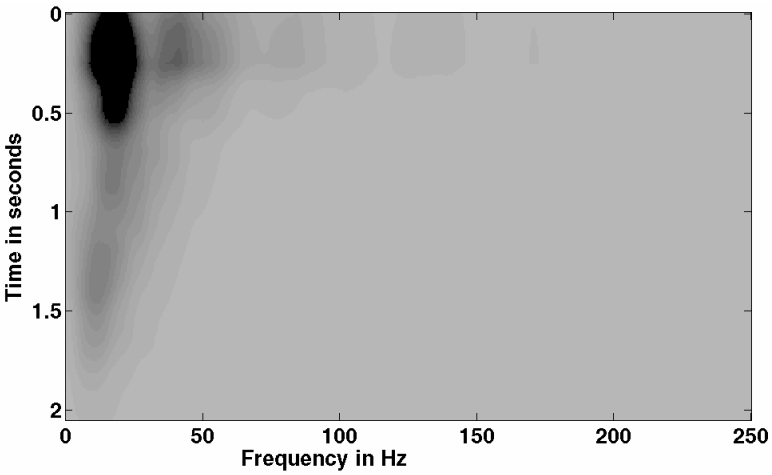


FIG. 6. The result of smoothing the Gabor magnitude spectrum of Figure 5 by convolution with a 2D boxcar of dimensions .5sec by 10 Hz. This estimates the Gabor magnitude spectrum of the propagating wavelet. The phase spectrum is calculated from this under the assumption of minimum phase.

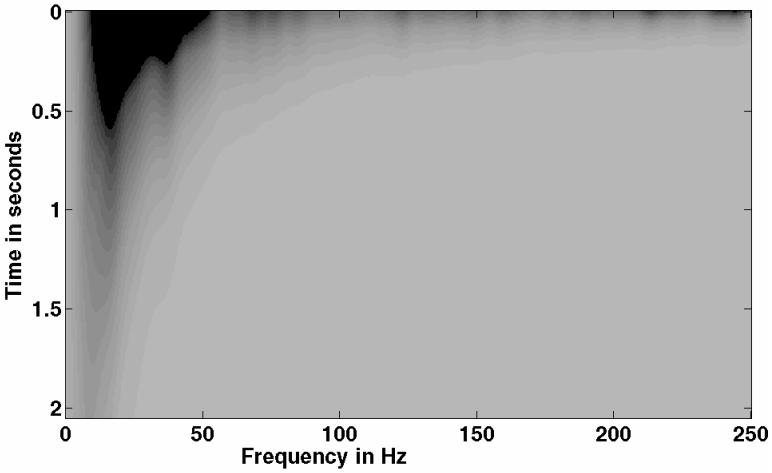


FIG. 7. Similar to Figure 6 except that the "hyperbolic smoothing" process, described in the text, was used. The phase spectrum is calculated from this under the assumption of minimum phase.

In Figures 8 and 9 we show the result of the spectral division of equation (19) using the propagating wavelet estimates of Figures 6 and 7 respectively. These are estimates of $G[r](n,f)$ and have been bandlimited along a hyperbolic trajectory between 1 and 2 seconds and along a vertical trajectory above 1 second. For comparison, Figure 10 shows the actual $G[r](n,f)$ without bandlimiting. We will see in a moment that the estimate of Figure 9 is superior to that of Figure 8. For further comparison, Figure 11 is the Gabor magnitude spectrum of a reflectivity estimate from Wiener's algorithm run on the same seismic signal. The stationary deconvolution was run between two passes of AGC (automatic gain correction) for amplitude control (if not it would be much worse). The design window for the Wiener algorithm was from .8 to 1.2 seconds. Figure 11 shows that there are still strong evidences of the attenuation process remaining in the Wiener result. Furthermore there is evidence of overwhitening above the design gate and underwhitening below the design gate.

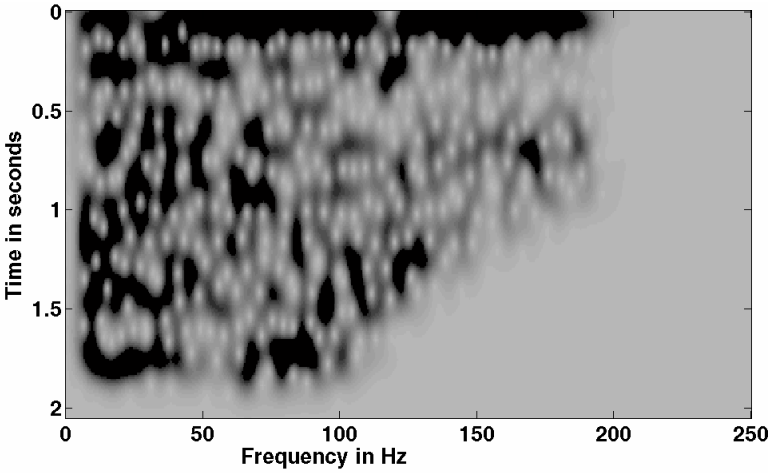


FIG. 8. The magnitudes of the estimate of the Gabor spectrum of the reflectivity achieved by dividing the input Gabor spectrum (Figure 5) by the estimate of the propagating wavelet achieved with boxcar smoothing (Figure 6) Compare with the answer in Figure 10.

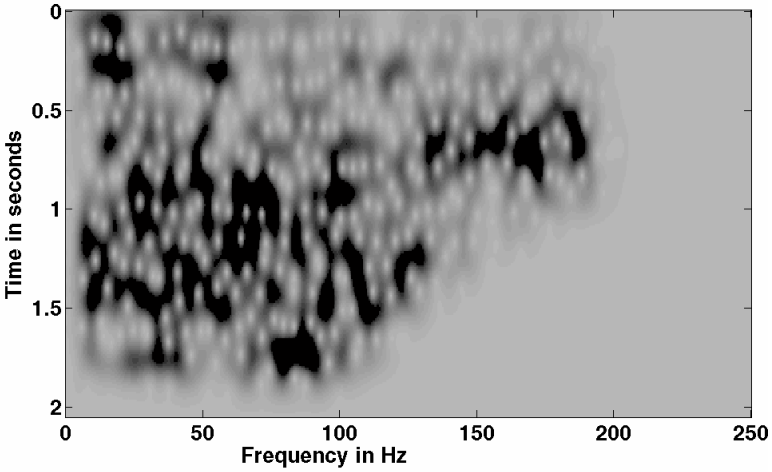


FIG. 9. The magnitudes of the estimate of the Gabor spectrum of the reflectivity achieved by dividing the input Gabor spectrum (Figure 5) by the estimate of the propagating wavelet achieved with hyperbolic smoothing (Figure 7) Compare with the answer in Figure 10.

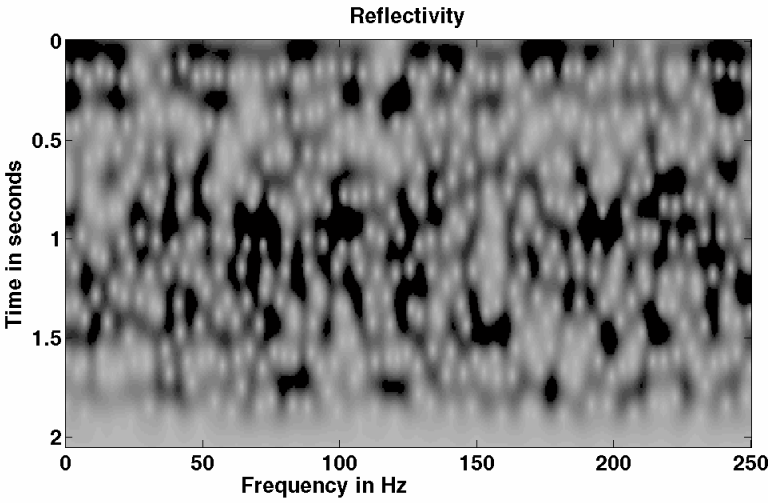


FIG. 10. The Gabor magnitude spectrum of the actual reflectivity used in the simulation.

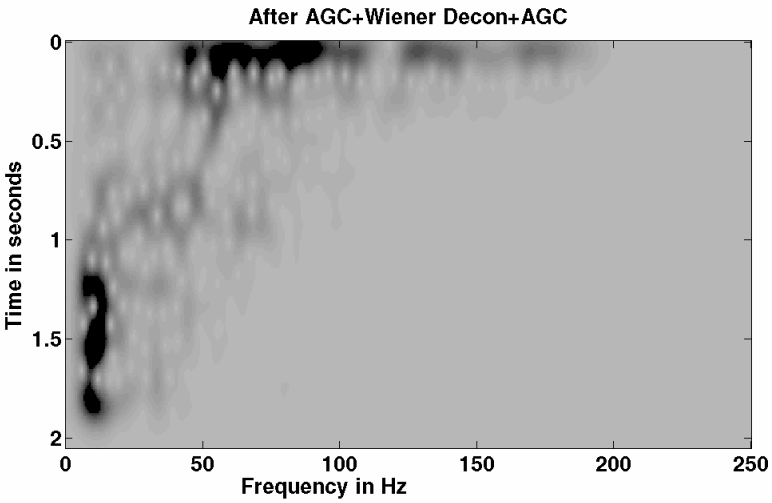


FIG. 11. The Gabor magnitude spectrum of the result of AGC → Wiener deconvolution → AGC run on the nonstationary signal of Figure 5. The Wiener design gate was .8 to 1.2 seconds. Note the overwhitening in the early times and the underwhitening in the later times.

Figures 12 and 13 show time-domain trace displays for the boxcar and hyperbolic smoothing processes respectively. Each display shows (from the bottom up) the input attenuated seismogram, the Wiener deconvolution result, the Gabor deconvolution result, and the bandlimited true reflectivity. (The reflectivity was given the same bandlimit as the Gabor spectra show in Figures 8 and 9). The only difference in these two figures is the Gabor result, all other traces are identical. Comparison of the two Gabor results shows the somewhat subtle improvement obtained by the hyperbolic smoothing process. Consider the large, isolated spike on the bandlimited reflectivity near 1.1 seconds. The result from boxcar smoothing shows this spike has been roughly equalized in amplitude with its neighbours. In contrast, the hyperbolic smoothing result has better preserved the exceptional amplitude of this spike.

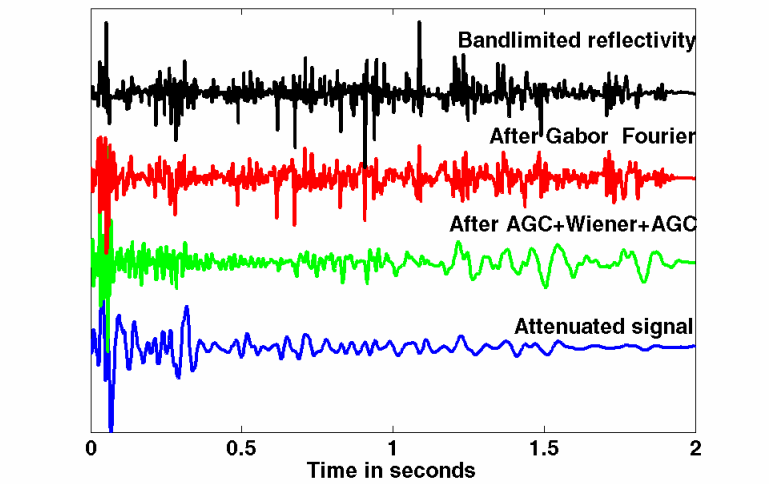


FIG. 12. A time-domain comparison of Gabor deconvolution and Wiener deconvolution for the synthetic nonstationary signal of Figure 5. The input signal is on the bottom and the desired answer is on the top. In this case, Gabor deconvolution used boxcar smoothing.

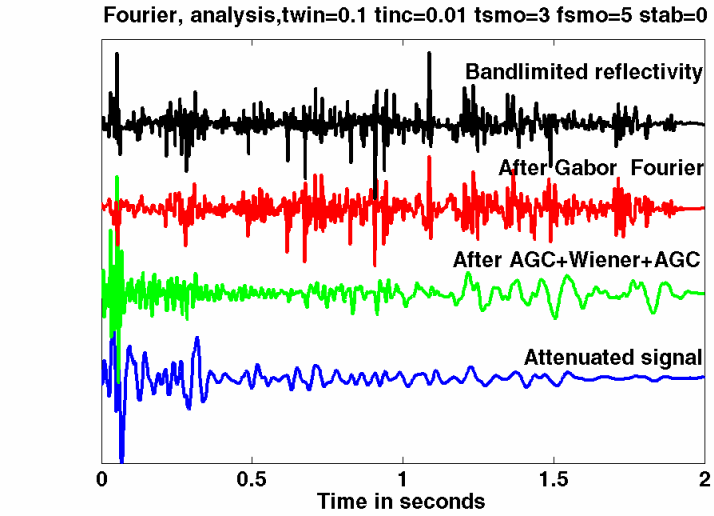


FIG. 13. A time-domain comparison of Gabor deconvolution and Wiener deconvolution for the synthetic nonstationary signal of Figure 5. The input signal is on the bottom and the desired answer is on the top. In this case, Gabor deconvolution used hyperbolic smoothing. In comparison with Figure 12, only the red signal (Gabor) is different.

The Wiener deconvolution result is also interesting in these figures. Clearly this deconvolution assignment was a task well outside the design assumptions of the Wiener algorithm. However, so also is the task posed by real seismic data. Close inspection of the Wiener result in the design gate (.8-1.2 seconds) shows that it is a reasonably good result there, with some resolution of the major events. However, even over this design gate, where the result should be 'optimal' there is clear evidence of spectral decay. At earlier times, there is evidence of overwhitening and also there is obvious underwhitening (poor resolution) at later times.

## Compact Gabor windows for Computational Efficiency

As an alternative to the Gaussian window suggested by Gabor, we have adopted the use of the window shown in Figure 14, which we call a Lamoureux window. The elementary Lamoureux window reaches unity at a single central point and is a polynomial on either flank. The flanks are chosen such that a point that is $\delta t$ from the beginning of the window has amplitude say $a$ and the point that is $\delta t$ from the central peak has amplitude $1-a$. This property allows an easy construction of a POU. A smooth window is found by choosing the unique odd polynomial of degree $2k+1$ which has its first $k$ derivatives equal to zero at $x=1$, and satisfies $p_k(1)=1$, then shift and glue together these pieces. For instance, such polynomials are given as

$$
\begin{aligned}
p_0(x) &= x \\
p_1(x) &= \left(3x - x^3\right)/2 \\
p_2(x) &= \left(15x - 10x^3 + 3x^5\right)/8 \\
p_3(x) &= \left(35x - 35x^3 + 21x^5 - 5x^7\right)/8 \\
&\vdots
\end{aligned}
\tag{20}
$$

and the corresponding Lamoureux window is given by

$$
wlam_k(x) = \begin{cases}
.5\left(1 + p_k(x)\right) & . \quad -1 < x < 1 \\
1 & . \quad 1 < x < L-2 \\
.5\left(1 + p_k(L - x - 1)\right) & . \quad L-2 < x < L \\
0 & . \quad \text{otherwise}
\end{cases}
$$

In comparison with a Gaussian window of the same width, the Lamoureux window is much more localized. (The width of the Gaussian is defined as its 1/e point.) In the context of the theory of the Gabor transform, if we choose $p = 1/2$, then we will actually window with the square root of the Lamoureux window in both analysis and synthesis. For this reason, a 4th order Lamoureux is a natural choice so that the window can approach zero as a 2nd order curve. In contrast, if we used the popular raised cosine window, then its square root approaches zero linearly and has discontinuous slope at its endpoints.

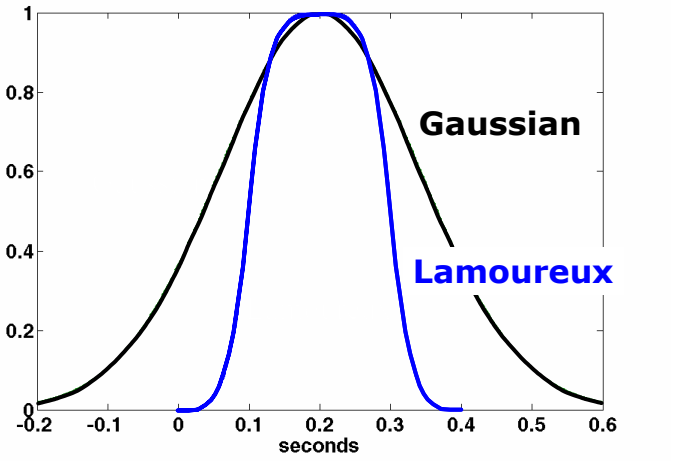


FIG. 14. A Gaussian window and a comparable Lamoureux window.

The use of compactly supported windows leads to great efficiencies in the Gabor transform. When using Gaussian windows, each windowed trace segment is as long as the original trace. So if $n$ windows are required and the trace is of length $N$, then the computation effort is proportional to $nN \log_2 N$. In contrast, if the compactly supported

window is of length $M$ where $M$ is typically $0.1N$ or less, then the effort is $nM \log_2 M \approx .1nN \log_2 N$. So we expect at least an order of magnitude less effort.

When using Gaussian windows, the window width and increment are essentially independent. In contrast, to form a simple POU with Lamoureux windows, the window spacing is determined by the width because the adjacent windows must have their peaks over the endpoints of the window in Figure 14. This is a minimally redundant POU (redundancy of 2). However, we can overcome this apparent obstruction to more redundant representations by simply shifting the minimally redundant partition by, say, 1/2 a window width and summing it to its originally. By this scheme, we can achieve any integral redundancy.

Figure 15 compares the Gabor transform using Gaussian windows to one using minimally redundant Lamoureux windows with the same nominal width. The Gaussian transform is very redundant (about 10 fold) and shows much higher resolution than the twice redundant Lamoureux window result. Nevertheless, both transforms can reconstruct the signal with high fidelity.

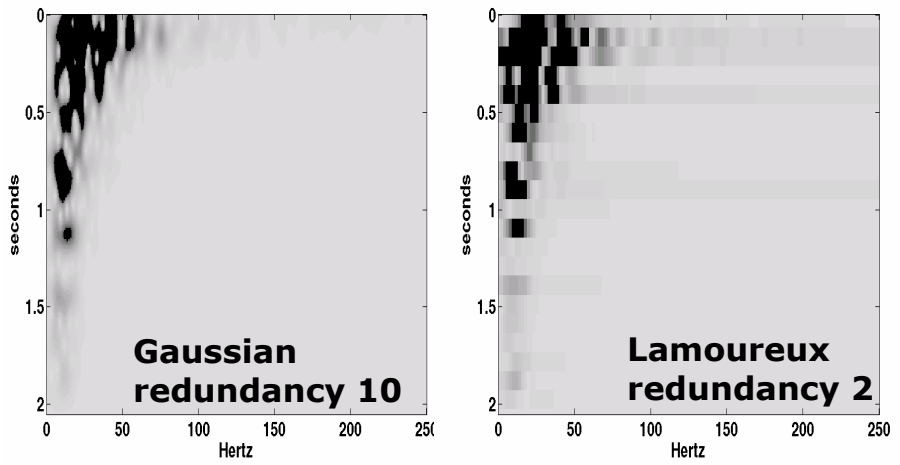


FIG. 15. A Gabor spectrum computed with a Gaussian window (left) and a Lamoureux window (right).

## REAL DATA COMPARISON

We now present an extensive comparison with real data. Most of the data processing was not done by us, but rather by Sensor Geophysical of Calgary, Alberta. Sensor has a reputation of be an excellent data processing shop that specializes in very high resolution images. Their standard processing flow, developed with a great deal of careful study and experimentation, is highly sophisticated, and we only comment on the steps required for comparison with Gabor deconvolution. Sensor's standard flow begins with true-amplitude gain recovery and continues with surface consistent Wiener deconvolution. This is followed by TVSW (time variant spectral whitening), CMP Stack, TVSW again, and FX noise reduction. TVSW is an interesting process (see Yilmaz, 1987) that does a strong, nonstationary spectral whitening that is effectively zero phase. The algorithm essentially uses a POU (partition of unity) in the Fourier domain to specify a set of narrow band filters that break a signal into a given number of frequency slices. These

filters are chosen to span the expected signal band and typically have a 5 or 10 Hz width. Since the filter's spectral shape is based on a POU, the frequency slices sum to recover the original signal. The TVSW process runs an AGC on each filter slice (in the time domain) and then sums them  This process has many virtues including rendering the amplitude spectrum essentially stationary, trace balancing, and coherent noise suppression.  Though not obvious, the results are very similar to running Gabor decon using a box car smoother and associating a zero phase spectrum with the $|\hat{w}_\alpha|$ estimate. For comparison, we have run two different Gabor deconvolution flows.  In what we will refer to as "pre stack Gabor" or PSG we substituted a Gabor deconvolution for the sequence "true-amplitude recovery, surface consistent deconvolution, and TVSW".  In this flow, TVSW and FX noise reduction remained in the post stack position.  In a second flow that we will call "pre and post stack Gabor" or PPSG we preceded the pre stack Gabor deconvolution with a radial trace filtering process (Henley, 20??) designed to attenuation coherent noise.  Then we also replaced the post stack TVSW with a post stack Gabor deconvolution.

Rarely in seismic data processing is a comparison of processes without secondary complications.  We have several here.  First, we found that we could recover a much broader signal band with Gabor deconvolution than was attempted with the standard processing.  The standard processing recovered signal out to about 125 Hz while we found that we could push this number to 160 Hz with Gabor deconvolution (this is 2 mil data so the Nyquist frequency is 250 Hz.).  At this time we do not know whether the standard flow could be pushed higher.  Another complication is that we chose to insert radial trace filtering in the PPSG flow because we hoped to demonstrate recovery of reflection signal at very low frequencies.  The radial trace filtering process is very effective at removing coherent source noise which might allow the low-frequency reflection signal to be recovered.  We are still assessing this point and will not report further on it here.

We have not done a great deal of parameter testing.  The Gabor algorithm has several important parameter choices including the selection of window type (Gaussian or Lamoureux), window width, window spacing, choice of "p" value in equation (5), smoother algorithm and smoother lengths, and white noise level.  For this case we chose Lamoureux windows of .2 seconds long, 50% window overlap, p=.75, hyperbolic smoother with a 10 Hz frequency boxcar, and a white noise level of .0001 .  It may well be that different parameter choices could provide superior results than shown here though we have found the algorithm to be quite robust with the parameters as given.

Our test line was provided courtesy of Husky Energy and was acquired over a sedimentary basin using a dynamite source.  Figures 16, 17, and 18 show a portion of our seismic line processed by the standard flow, the PSG flow, and the PPSG flow respectively.  The bandwidth of both Gabor results is visibly greater, even at this display scale (and the residual random noise is somewhat greater, as well). The overall amplitudes of the PPSG results are more uniform in time than either of the other results, indicating possibly too much amplitude levelling caused by two passes of the Gabor algorithm.  We conjecture that this AGC effect could be lessened with alternative parameters in the post stack Gabor.

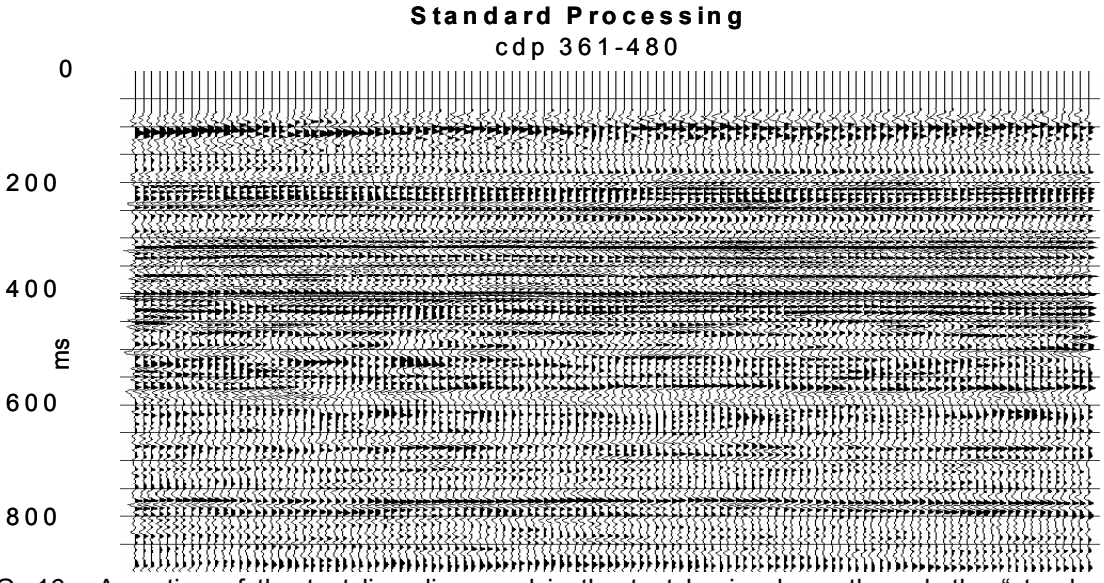


FIG. 16. A portion of the test line discussed in the text having been through the "standard" processing flow. Relevant steps in the standard flow are: "gain→ surface-consistent Wiener deconvolution → TVSW", in the prestack flow and "TVSW→ fx noise reduction" in the post stack processing.

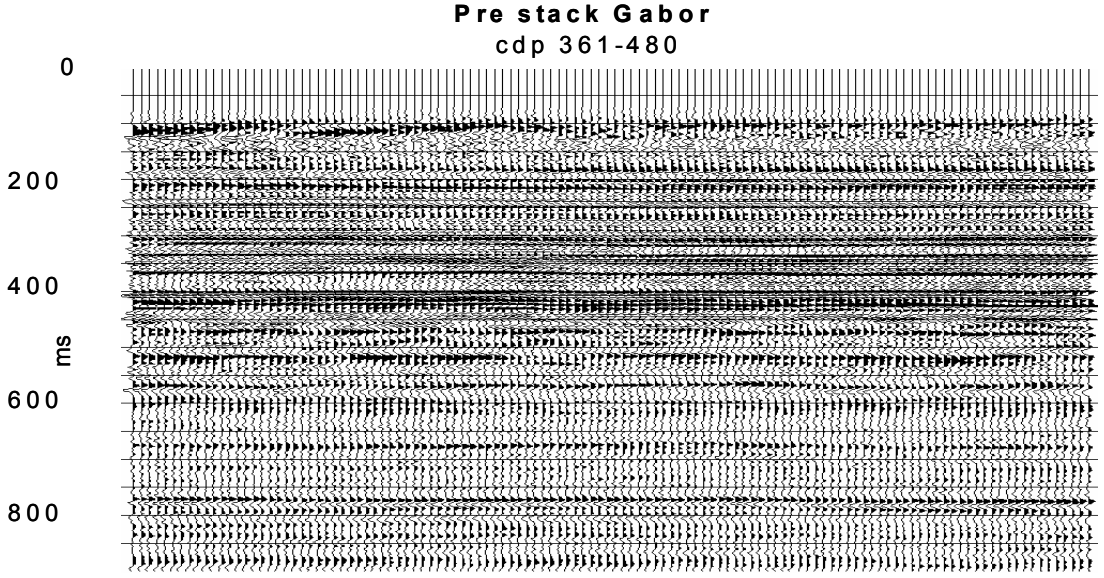


FIG. 17. The same portion of the line as shown in Figure 16 as processed through the "pre stack Gabor" flow. This flow replaces "gain→ surface-consistent Wiener deconvolution → TVSW" in the prestack portion of the standard flow with Gabor deconvolution.

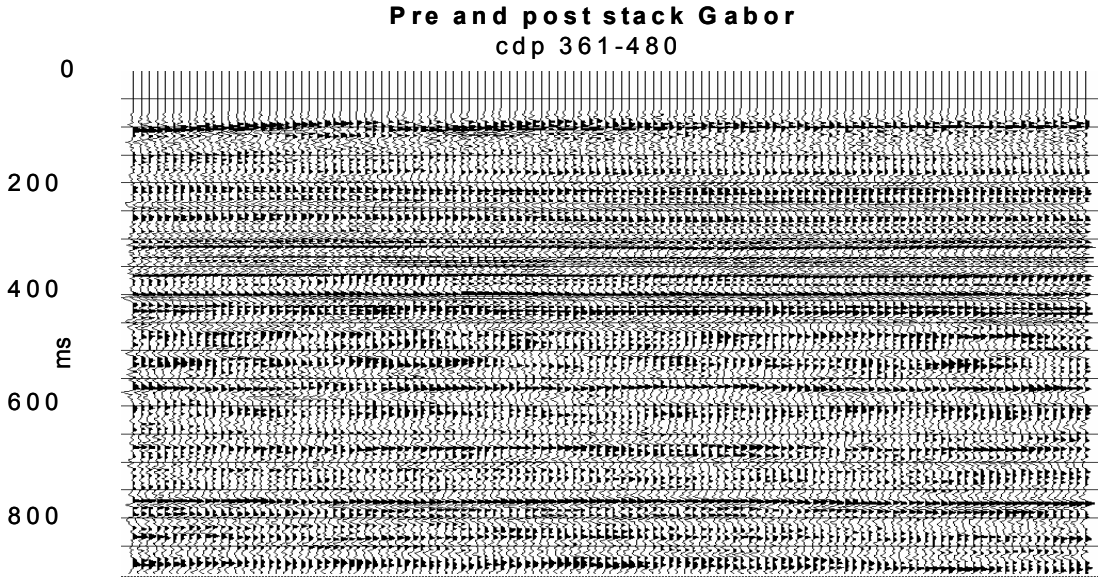


FIG. 18. The same portion of the line as shown in Figure 16 as processed through the "pre and post stack Gabor" flow. This flow replaces "gain→ surface-consistent Wiener deconvolution → TVSW" in the prestack portion of the standard flow with Gabor deconvolution and TVSW in the post stack flow with Gabor deconvolution.

Figures 19, 20, and 21 show cdp's 481-600 for standard, PSG and PPSG flows, respectively. The highlighted feature is a small gap in surface coverage, which seems to be effectively 'healed' by both versions of Gabor processing. The single-channel nature of our Gabor code means that it is also rapidly adapting spatially as well as temporally. Surface consistent deconvolution, in contrast, is constrained to change quite slowly. Amplitudes of most events appear more laterally consistent on both Gabor results than on the standard version, even though the Gabor results are greater in bandwidth. Once again, the PPSG results look a bit 'bland', as if the amplitudes have been levelled too much. However, the PPSG result appears to have more sharply defined events than the other two sections.

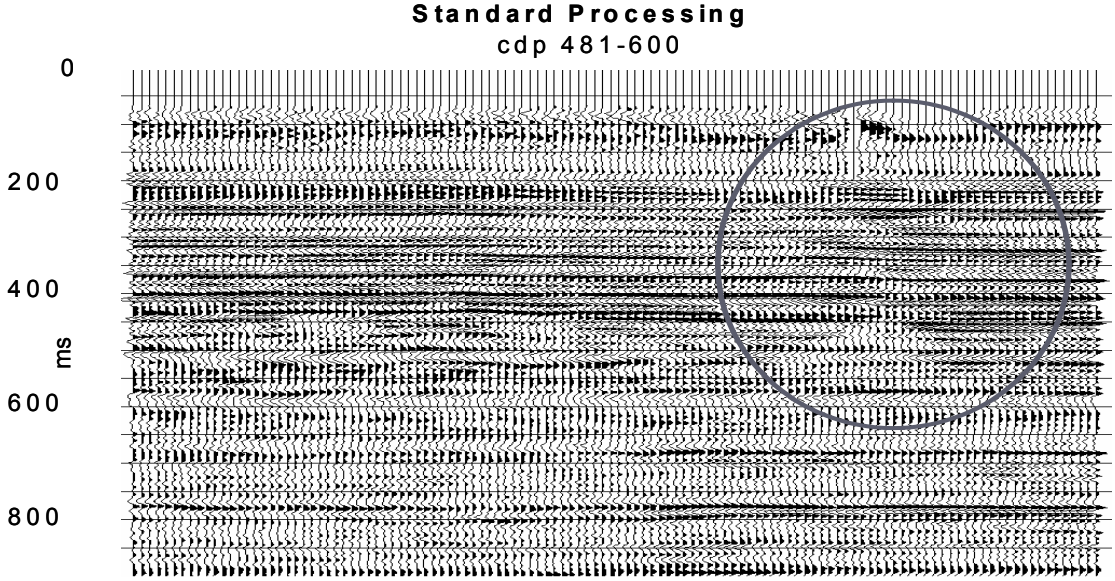


FIG. 19. Another portion of the test line as processed through the standard flow. The circle indicates an amplitude and character anomaly that arises due to a coverage gap.

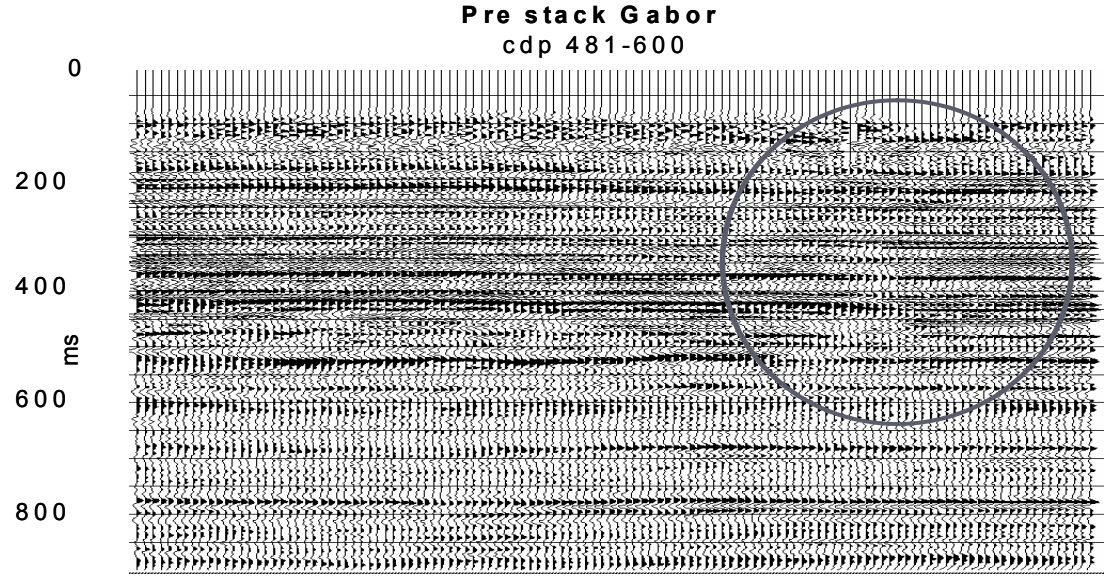


FIG. 20. The same portion of the test line as shown in Figure 19 but processed through the prestack Gabor flow. The circle indicates an amplitude and character anomaly that arises due to a coverage gap. The anomaly is much less after Gabor deconvolution.

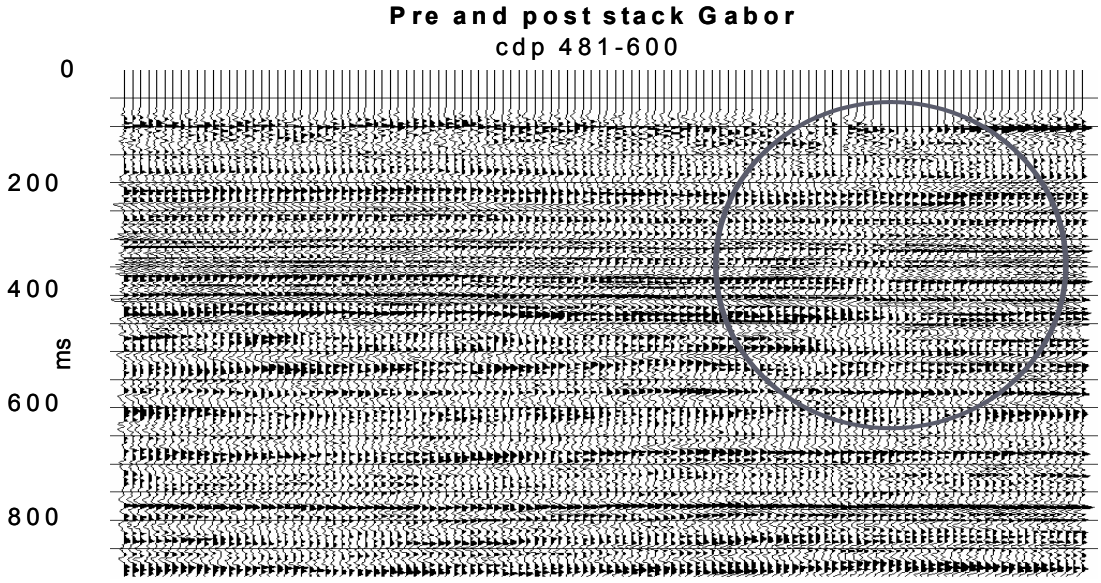


FIG. 21. The same portion of the test line as shown in Figure 19 but processed through the pre and post stack Gabor flow. The circle indicates an amplitude and character anomaly that arises due to a coverage gap. The anomaly is much less after Gabor deconvolution.

The outlined feature in Figures 22, 23, and 24 is likely a geological anomaly. It is more precisely delineated by the increased bandwidth of the Gabor results. In addition to greater bandwidth, the Gabor results appear to have better continuity as well.

Figure 25 is a single group of cdp's, 361-420, in which the leftmost group of 20 traces is from the PSG result, the centre 20 traces from the standard result, and the rightmost group from the PPSG flow. As can be seen, in spite of fairly different appearances as groups of traces, the trace events actually tie almost loop for loop. What causes their different appearance is apparently the bandwidth, not only the high frequencies, but the low frequencies as well. As a good example, the event sequence between 200 and 250 ms looks considerably different (especially the triplet of loops just below 200 ms on the standard section).

Figures 26, 27, and 28 show cdp's 661-720. The two Gabor results have visibly greater bandwidth but without any sacrifice in continuity. For example, not the doublet event whose bottom loop in right on 350 ms on the right hand side of the PPSG result. This event is continuous across most of the section but is really not discernable on the standard result.

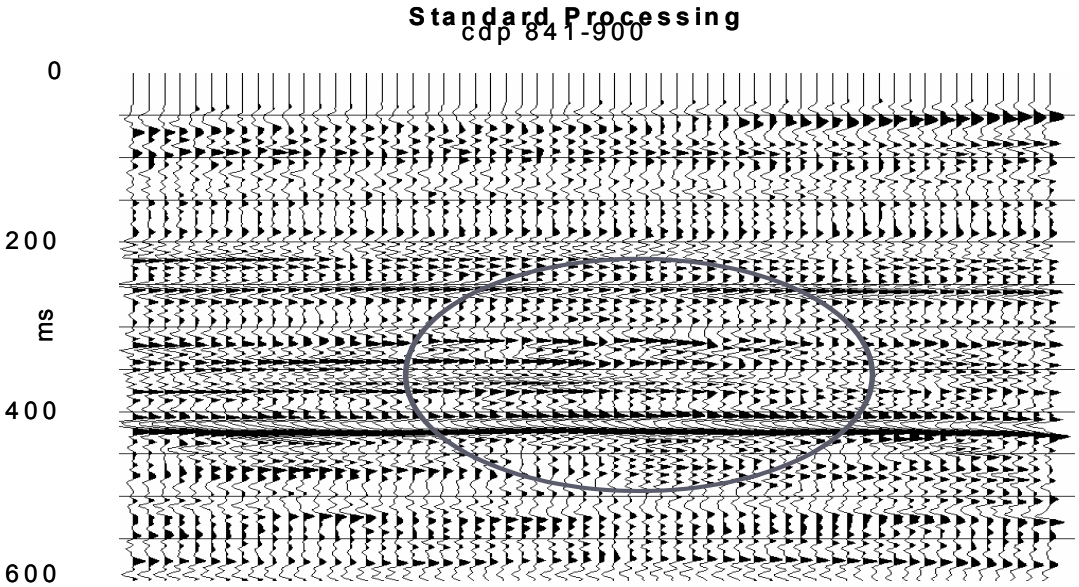


FIG. 22. Another portion of the test seismic line as processed with the standard flow. The circle highlights an apparent geological structure.

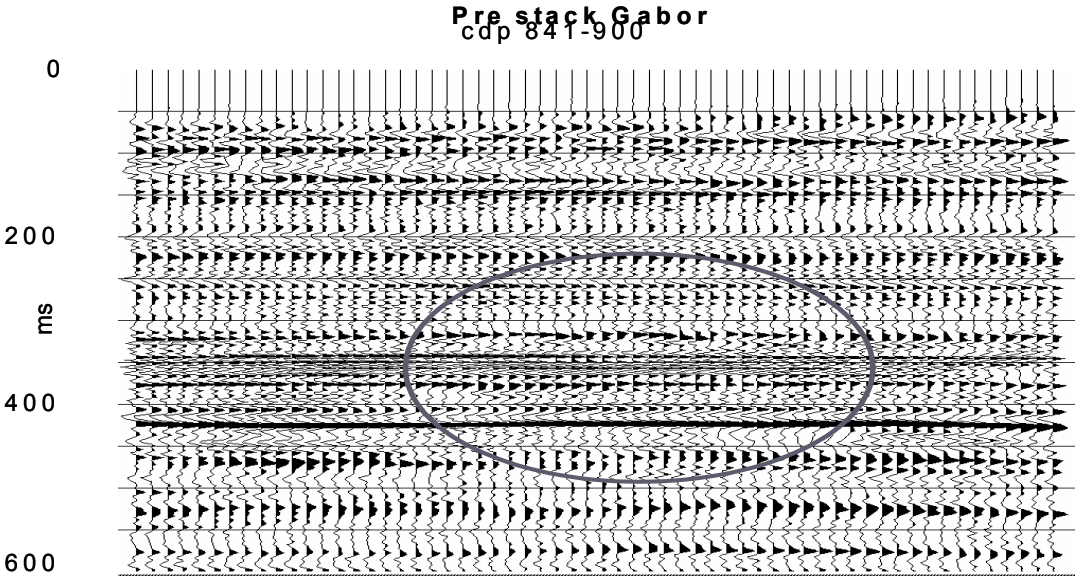


FIG. 23. The same portion of the test seismic line as Figure 22 but processed with the prestack Gabor flow. The circle highlights an apparent geological structure. Note the increased temporal resolution.

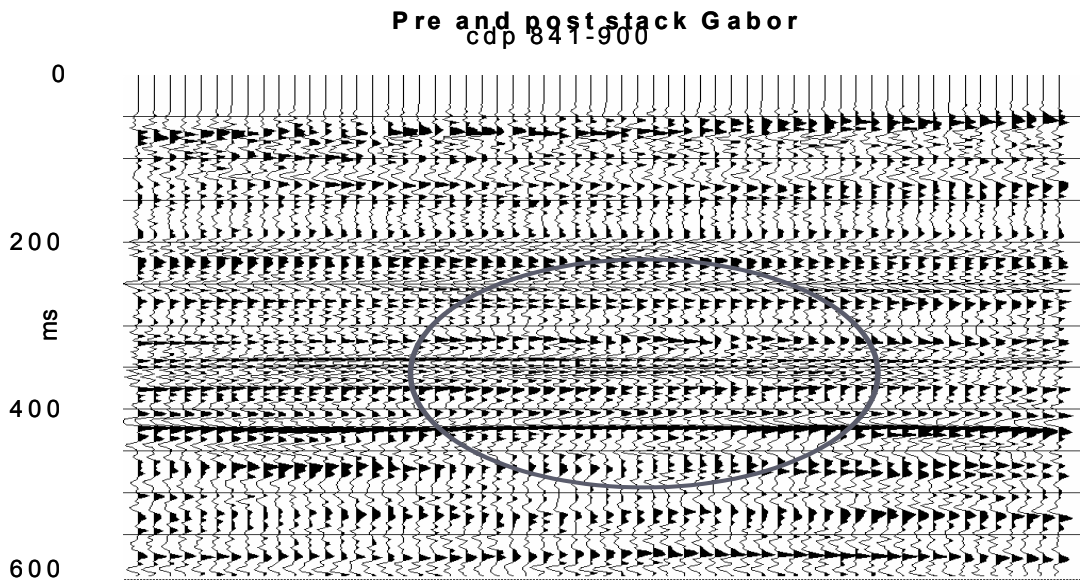


FIG. 24. The same portion of the test seismic line as Figure 22 but processed with the pre and post stack Gabor flow. The circle highlights an apparent geological structure. Note the increased temporal resolution.

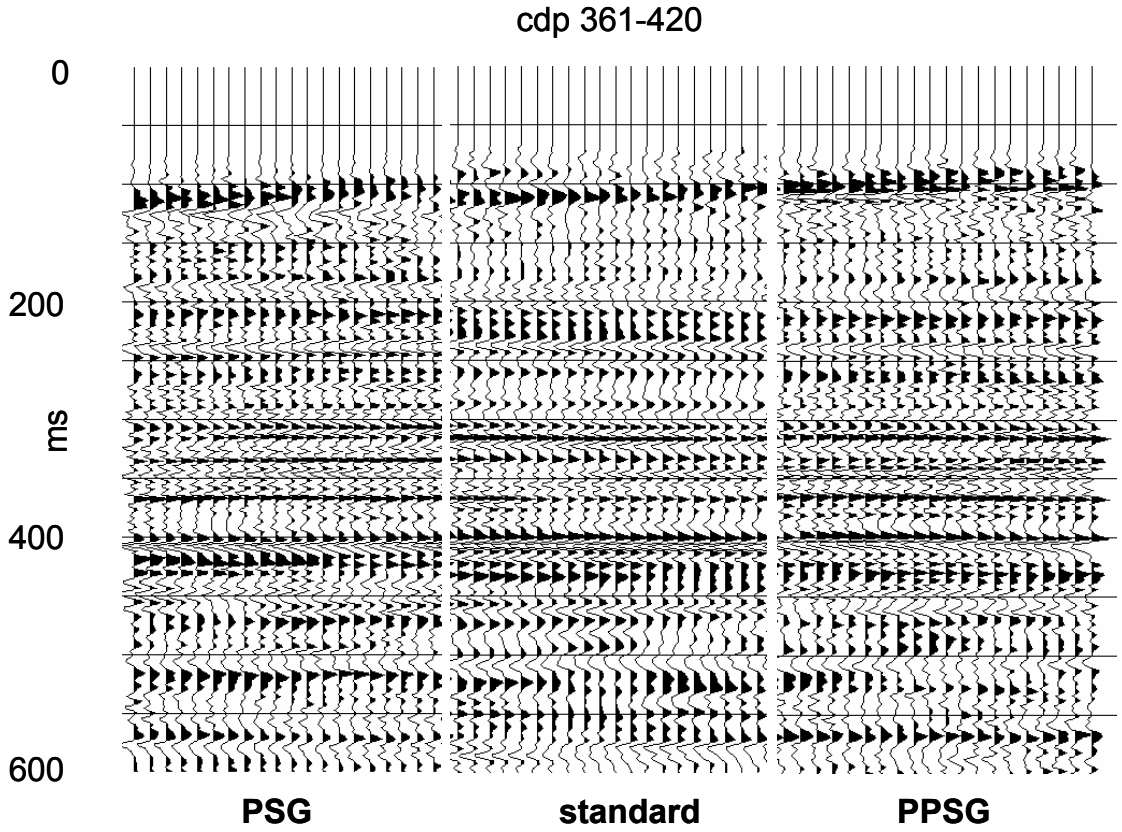


FIG. 25. A portion of the test line with three different panels of processing. The panels represent consecutive traces and are not the same traces repeated.

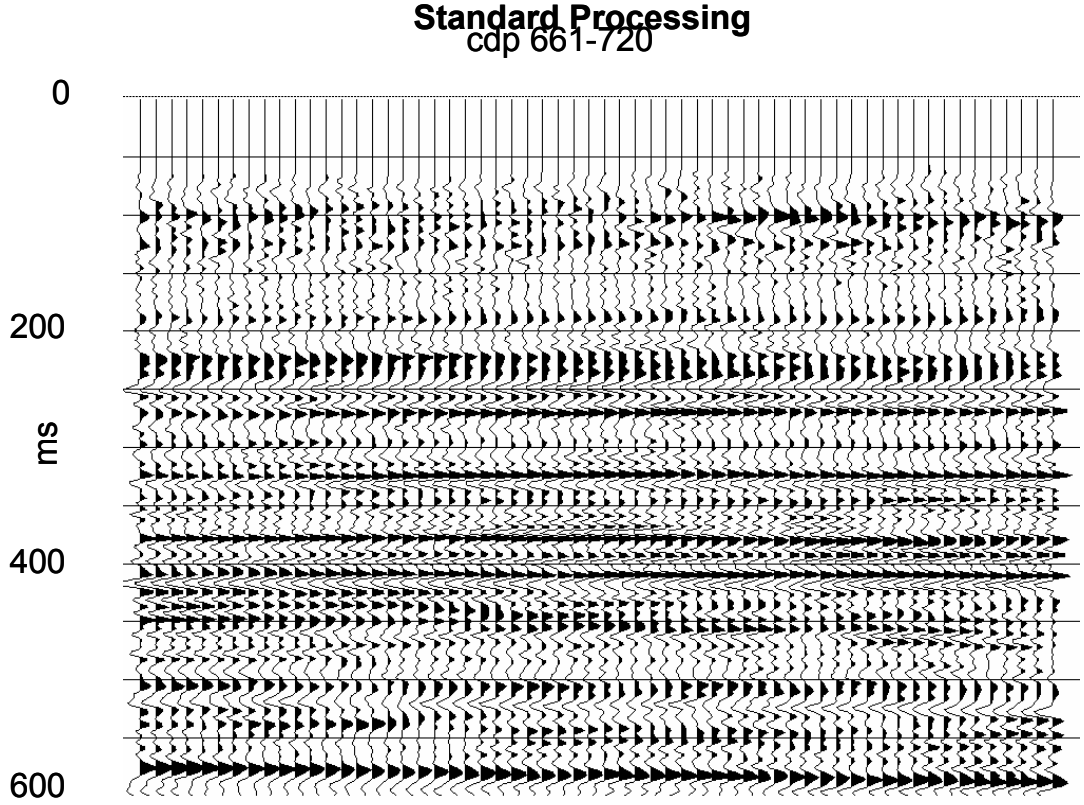


FIG. 26. Another portion of the test line imaged by the standard processing flow.

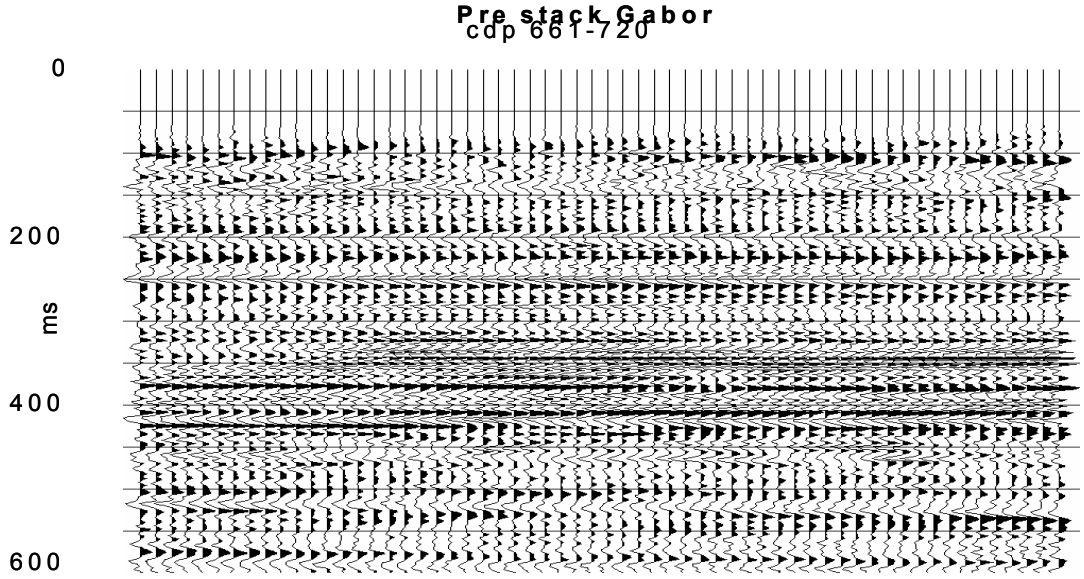


FIG. 27. The same portion of the test line as in Figure 26 but processed by the prestack Gabor flow. Note the much higher temporal resolution. Especially near 350 ms.

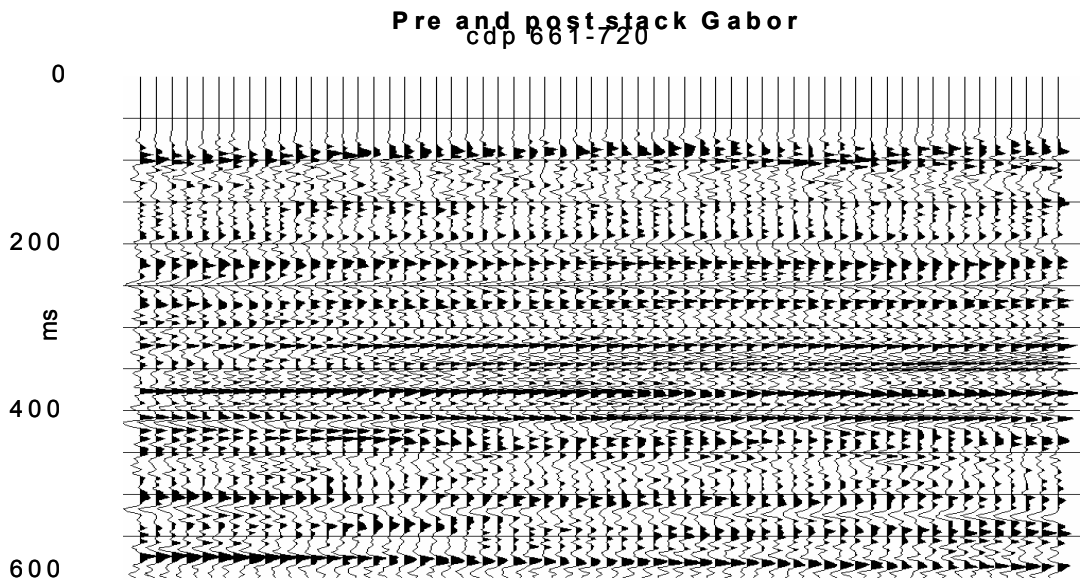


FIG. 28. The same portion of the test line as in Figure 26 but processed by the pre and post stack Gabor flow. Note the much higher temporal resolution. Especially near 350 ms.

## A VSP experiment

With the evidence in the previous section, it might be concluded that Gabor deconvolution gives an incrementally higher resolution image than the standard flow. However, the validity of that image may still be doubted. As an additional test, we have conducted an analysis of a unique dataset provided to us by Encana. This consisted of 5 surface vibration points at various offsets from a well with 3C receivers in a surface spread and in a VSP configuration downhole. Each sweep was recorded simultaneously into both receiver arrays.

We followed standard practice to perform wavefield separation on the VSP in order to directly estimate the downgoing waveforms. We then estimated the apparent waveforms in the surface data using three different methods: Gabor deconvolution, time-domain Wiener deconvolution, and frequency-domain Wiener deconvolution. The Wiener methods were done on gained data in temporally short analysis windows that were moved down the traces in increments. These were the same windows used by Gabor deconvolution. The Gabor estimates were made on raw data. At this point, we had wavelets observed in the VSP as a function of depth versus those observed on the surface versus two-way traveltime. These are not directly comparable since the VSP-observed wavelets have only travelled half of the travelpath of the surface-observed wavelets. Therefore we estimated Q values from the VSP using spectral ratios and developed a forward Q filter to apply to the VSP-observed wavelets to simulate travel along the two-way path to the surface.

Figure 12 typifies our results. In all cases, we observed that the Gabor estimates agreed quite closely with the direct VSP observations. While the Wiener estimates were close to one another, they differed significantly in amplitude and phase from the other two, especially in the deeper window. While the amplitude effects are simply attributable to the Wiener methods requiring gained data, the phase differences are not. In any case,

without overly criticizing the Wiener results, it is pleasing that the Gabor results agree so closely with the VSP observations.

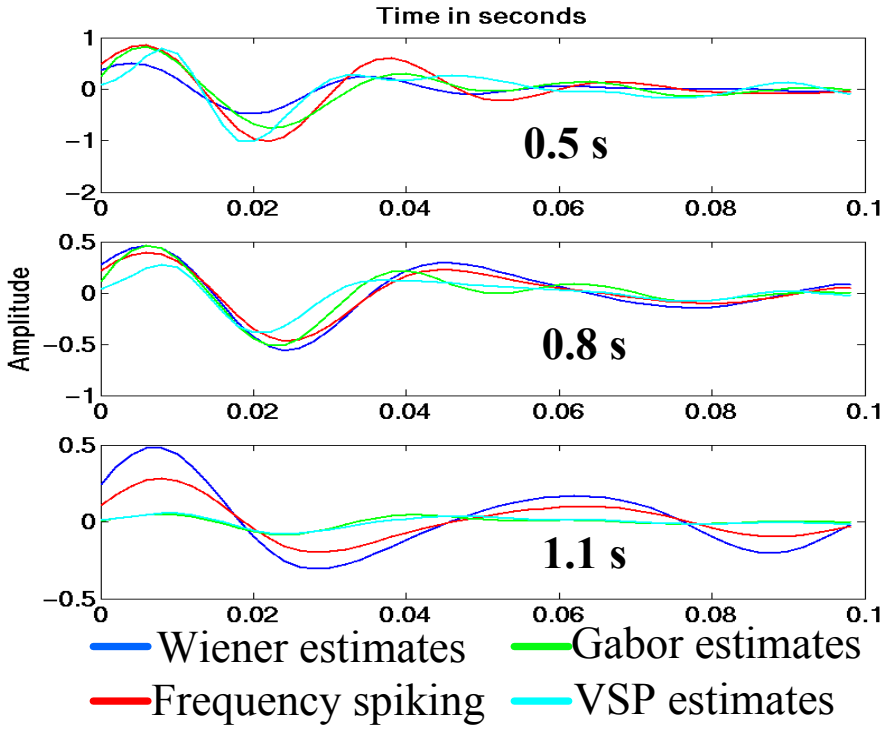


FIG. 29. Wavelet estimates from another test dataset. The estimates labelled "Wiener", "Frequency spiking", and "Gabor" are all from the indicated deconvolution algorithms on surface recorded data. The VSP estimates are direct measurements in a borehole by standard VSP wavefield separation.

One final note, it surprised us, given the Vibroseis source, to find that the VSP observations did not require a phase rotation to agree with the minimum phase assumptions of the deconvolution algorithms. This direct observation of apparently (or perhaps "effectively") minimum-phase wavelets from a VSP experiment has been made by others (notably Larry Mewhort of Husky Energy) and deserves more attention than we can give it here.

## CONCLUSIONS

We have presented a new deconvolution algorithm that extends the standard Wiener algorithm to directly address nonstationarity of the propagating wavelet. Called Gabor deconvolution because it utilizes the Gabor transform, a nonstationary extension of the Fourier transform, our method is based on an extension of the standard convolutional model of the seismic trace to a nonstationary convolution. This new nonstationary convolutional model explicitly accounts for the waveform as emitted by the source and for its subsequent continuous modification by attenuation processes as it propagates. This model incorporates greater physical realism than the standard convolutional model.

We presented a simple mathematical framework for the Gabor transform of a discretely sampled, bandlimited signal. Our discrete Gabor transform recovers a signal

exactly by using a set of windows that form a partition of unity. Furthermore, we can apply windows both at the analysis and at the synthesis stages.

Mathematically, our algorithm succeeds because the Gabor transform approximately renders our nonstationary convolutional model into a product of factors. This is similar to the fundamental fact that the Fourier transform factorizes a convolution integral except that the Fourier factorization is exact while the Gabor factorization is not. Our model predicts that the Gabor transform of a seismic signal is the product of (1) the Fourier transform of the waveform emitted by the source, (2) the time-frequency attenuation function, and (3) the Gabor transform of the reflectivity. Since it is the latter element that we desire to estimate, this suggests the possibility of estimating the product of the first two, which we call the propagating wavelet, and dividing by them in the Gabor domain.

Invoking the white reflectivity assumption, we argue that it is possible to develop a smoothing process that, when applied to the magnitude of the Gabor spectrum of the seismic trace, will eliminate all vestiges of the reflectivity and leave only the effects of source signature and attenuation function. We then recover the phase of these two factors by invoking the standard minimum phase assumption. Given this estimate of the propagating wavelet, Gabor deconvolution is completed by a spectral division in the Gabor domain followed by an inverse Gabor transform.

We demonstrated our algorithm on a nonstationary synthetic signal where it is clearly superior to Wiener deconvolution. We used a real data example to compare Gabor deconvolution to the common practice of surface-consistent Wiener deconvolution followed by TVSW (time variant spectral whitening). Though less compelling than with a synthetic there still seems to be a significant advantage in resolution from using Gabor deconvolution.

Finally we showed an example with real data where a seismic wavefield was recorded simultaneously into a VSP and a surface spread. Wavelets were estimated directly from the VSP using standard wavefield separation methods and compared with those estimated on the surface by Gabor deconvolution. Good agreement was found.

## ACKNOWLEDGEMENTS

We thank the sponsors of both the CREWES and POTSI projects. These include a variety of industrial and government sources. We are especially grateful to Encana and Husky Energy for providing the data examples and to Sensor Geophysical for providing data processing. Some of this work was accomplished while some of the authors were in-residence at the Banff International Research Station (BIRS) for the mathematical sciences.

## REFERENCES

Bastiaans, M. J., 1980, Gabor's expansion of a signal into Gaussian elementary signals: Proceedings of the IEEE, **68**, 538-539.

Feichtinger, H. G., and Strohmer, T., 1998, Gabor analysis and algorithms: Theory and applications: Birkhauser, ISBN 0-8176-3959-4.

Gabor, D., 1946, Theory of communication: J. IEEE (London), **93(III)**, 429-457.

Grochenig, K., 2001, Foundations of time-frequency analysis: Brikhauser, ISBN 0-8176-4022-3.

Kjartansson, E., 1979, Constant Q-wave propagation and attenuation: Journal of Geophysical Research, **84**, 4737-4748.

Kohn, J. J., and Nirenberg, L., 1965, An algebra of pseudodifferential operators: Comm. Pure and Appl. Math. **18**, 269-305.

Margrave, G. F., 1998, Theory of nonstationary linear filtering in the Fourier domain with application to time-variant filtering: Geophysics, **63**, 244-259.

Margrave, G. F., and Lamoureux, M. P., 2002, Gabor Deconvolution: 2002 CSEG Annual Convention, Calgary, AB.

Peacock, K. L., and Treitel, S., 1969, Predictive deconvolution: Theory and practice: Geophysics, **34**,155-169.

Robinson, E. A. and Treitel, S., 1967, Principles of digital Wiener filtering: Geophys. Prosp., **15**, 311-333.

Schoepp, A. R., and Margrave, G. F., 1998, Improving seismic resolution with nonstationary deconvolution: 68th Annual SEG meeting, New Orleans, La.

Yilmaz, O., 1987, Seismic Data Processing: Soc. of Expl. Geophys., 526.