

Seismic deconvolution: assumptions, concerns and convictions

Peter W. Cary, Sensor Geophysical Ltd.



Summary

Early in our careers as geophysicists, most of us took at least one course on seismic signal analysis where we were taught that standard Wiener deconvolution converts the minimum-phase source wavelet in our seismic data to a wavelet with a phase spectrum that is zero and an amplitude spectrum that is broad and flat. If the seismic source is not minimum phase, such as for Vibroseis data, we need to convert the source wavelet to minimum-phase before deconvolution to ensure that our output is zero phase. It all seemed so straightforward.

Once we started to gain experience in the industry, we found that the standard wavelet processing flow is a little more complicated than simply running Wiener deconvolution, as we learned in school. Some extra steps, like instrument dephasing or power-line noise removal, are run first, and we use surface-consistent deconvolution with lots of different components, not trace-by-trace decon. After that we run zero-phase deconvolution on a trace-by-trace basis, in order to whiten the amplitude spectrum, even though it is already supposed to be white. We may even do some more whitening after stack. There are a lot of variations on exactly how these steps are run, how many times each one is run, and what is done between each step, but we learned that deconvolution in general is still fairly straightforward. We still basically just pick an operator length, a prewhitening level, a deconvolution design gate, and submit the job. What could possibly go wrong?

Normally, we are not aware of deconvolution causing many problems. However, if there is one problem with deconvolution that continues to rear its ugly head, it is difficulties with phase. The imperfections in our wavelet processing can begin to appear when we compare our stacked data to well log-based synthetic seismograms. Although the match is often fairly good, it does not always turn out that way. Sometimes the stacked section comes out disconcertingly far from zero-phase, and once we look closer, we see that some parts of the seismic trace often fit the synthetics better than others. On these occasions, however, we can usually explain away the problem in terms of a bad section of the well log, or a polarity flip in the field, or noise, or multiples, or whatever. After all, we do not have time to track down every problem we encounter since there is always more data waiting to be processed or interpreted.

But more problems often crop up when we compare our recently acquired data to older vintages of data acquired in the same location. There is often some discrepancy in the phase and amplitude of reflectors that we have to alter, or hide, or explain away. This is not always the case—sometimes it actually comes out right, but phase-tying seismic to synthetics, and seismic to seismic, is an unfortunate part of our everyday geophysical lives that seems to have been forced upon us by some kind of failure in our wavelet processing to produce better results.

The severity of this phase infidelity problem depends at least partly on the beholder, and also, of course, on the character of the exploration target. If the object of interpretation is purely structural, details of phase will obviously not be as big a concern as for a subtle stratigraphic trap. As far as people's attitudes towards deconvolution go, there are some people who are amazingly content with the output of standard processing, and appear to believe that the output of deconvolution and the other wavelet processing steps is entirely trustworthy. At the opposite end of the spectrum, there are some who hold a strong distrust of the outcome from any deconvolution process. It seems that the majority, however, are noticeably cautious around deconvolution—they admit that the results are never completely right, but they believe (or hope) that the results are not grossly wrong either, at least after the data have been checked against synthetic seismograms.

Since a lot of geophysics operates in an openly imprecise state, where error bars are pretty much thrown to the wind, we tend not to let things like phase errors after deconvolution bother us too much. After all, interpreters ultimately have to trust the wavelet processing on their seismic data if the sections are to be interpreted and well locations are to be picked. But it is disconcerting, nonetheless, that we seem to be forced into this kind of submissive state by deconvolution. That is, we know that deconvolution is not reliable enough to produce a zero-phase output all the time, but we continue to use it, and obsequiously accept its flaws without complaint whenever we correct the results with phase rotation afterwards.

Is this a rhetorical overstatement of the real situation? Judging from the amount of time that seismic processors put into phase matching intersecting 2-D seismic lines, and overlapping 3-D seismic surveys, it appears to be simply a statement of the truth. Since Wiener deconvolution tries to estimate stationary wavelets from noisy, nonstationary seismic data, we should not be too surprised to see some random variation in the phase of our deconvolved output. What we observe, however, is not random. I have heard some people dismiss this problem in an off-handed manner by saying that it would go away if the data were “properly” processed, whatever that means. However, the situation is usually not that simple.

It is much more likely that the fluctuations in phase that we observe after deconvolution (the type of errors that we try to fix with constant phase-rotations afterwards) are caused by systematic errors in estimation that are brought about by a combination of the character of the noise in the data and some of the details in our processing flows. It is very easy to show, for example, that the phase of the data can be altered simply by changing the range of offsets in the deconvolution design gate, or by running a noise attenuation step in the flow before deconvolution. This variation in phase is caused by a relatively small systematic change in the estimation of the low-frequency part of the amplitude spectrum.

Since we assume a minimum phase source, we construct the phase spectrum from the log-amplitude spectrum via a Hilbert transform. However, knowledge of the entire log-amplitude spectrum is required for accurate phase estimation, even outside the signal bandwidth. Unfortunately the wavelet phase is most sensitive to the low-frequency part of the spectrum, which is the most difficult part to measure. Therefore, an error in the low-frequency part of the amplitude spectrum, which for land data is often highly contaminated with source-generated noise, can translate directly into an error in the phase of the data after deconvolution. This is a tough problem to solve. It is easy to come up with remedies, but none of them is fail-safe. This problem alone can make us realise how difficult it is to attain complete phase control.

Despite our knowledge of these types of phase errors from deconvolution, we still do not worry very much because we believe that the phase errors do not change laterally down a seismic line. As long as that is true, a single phase operator can straighten out the problem so that changes in wavelet character along a horizon can still be reliably interpreted as being due to changes in geology, not due to errors in the deconvolution. This assumption of lateral invariance is what allows us to sleep at night. Unfortunately, some recent studies indicate that our blissful state could sometimes be due to ignorance.

In addition to these problems, there are a host of other deconvolution issues that affect our day-to-day lives. Convictions about deconvolution are sometimes so strong that these types of issues are capable of transforming normally docile geophysicists into emotionally-charged zealots. Do we need to bother dephasing our data before deconvolution? Is there any point to all of these components in our surface-consistent deconvolution programs? Can any quality control product hope to tell us whether deconvolution has worked correctly or not? Do multiple passes of deconvolution, or noise attenuation before deconvolution, ever gain anything? Why are we whitening the amplitude spectrum when the reflectivity's spectrum is blue? Doesn't time-variant spectral whitening whiten the data too much? Are we ever able to process Vibroseis data so that it comes out zero phase? Shouldn't we always be inverse-Q filtering our data?

Some of these issues are so unimportant that it is ridiculous to argue about them. Others, however, are so important that it is ridiculous not to argue about them. Deconvolution can be such a touchy subject that even just the mention of some of these issues in this presentation is sure to generate controversy.