

Basic numerics in the nonlinear inversion of primaries: simultaneous imaging and inversion II

Kristopher A. Innanen, Arthur B. Weglein and Tad J. Ulrych

Abstract

The coupling of the tasks of imaging and inverting seismic primaries discussed in Innanen et al. (2004) (see also Innanen, 2003 and Innanen and Weglein, 2003), while from a practical viewpoint running contrary to the strategy of task-separation (e.g. Weglein et al., 2003), is valuable as a tool for the understanding of the functioning of the series as a whole. In this paper we numerically implement the formula for simultaneous imaging and inversion.

We begin by using the numerical implementation to discuss two basic strategies. (1) We stabilize the addition of high-order terms with a weighted cutoff of high-frequency/high-wavenumber portions of the series, a strategy that bears a strong resemblance to the so-called truncated singular value decomposition inverse methods. This is found to be necessary because of the reliance of the n 'th term in the sub-series' on the n 'th derivative of the input. (2) We demonstrate a candidate approach to the compensation for bandlimited input (i.e., missing low and high frequencies), using the "gap-filling", or spectral extrapolation techniques that exist in the literature for bandlimited impedance inversion. Finally, we use the analogy of the Taylor's series expansion of an exponential function to predict and explain the close relationship between (a) the order of the imaged/inverted model reconstruction (i.e. the number of terms used in its computation), and (b) the frequency content of the reconstruction. This relationship, simply put, says that a low-order truncation of the series corresponds to a low-frequency (or smooth) reconstruction. Adding more terms adds higher frequencies; this suggests the existence of a trade-off between resolution and truncation order that may be critical in practical computation.

Simple acoustic 1D normal incidence configurations for the numeric examples are used in order to make simple comments on the broad nature of various subseries, and the computation of truncated versions thereof. We expect these comments, although based on simplistic examples, to apply in their essence (although no doubt with some variation in detail) to more complicated instances of inverse scattering series algorithms (i.e., multidimensional, elastic generalizations), because of their reliance on the same basic math-physics formalism.

1 Introduction: Models and Numerical Issues

In Innanen et al. (2004), we investigated the derivation and properties of the quantity

$$\alpha_{SII}^{(n)}(z) = K_n \frac{d^n}{dz^n} \left(\int_0^z \alpha_1(z') dz' \right)^n, \quad (1)$$

where $\alpha_1(z)$ is the linear component of $\alpha(z) = 1 - c_0^2/c^2(z)$ for an acoustic, 1D normal incidence seismic problem, and

$$K_{n+1} = \frac{(-1)^n}{n+1} \left(\frac{1}{2^n} \right)^2 \left[\sum_{k=0}^n \frac{1}{k!(n-k)!} \right]. \quad (2)$$

In this section, we begin the investigation into the numerical computation of

$$\alpha_{SII}(z) = \sum_{n=1}^N \alpha_{SII}^{(n)}(z). \quad (3)$$

1.1 Synthetic 1D Models, Data and Born Approximations

In this section we use some simple Earth models to exercise the numerical computation of equation (3). We begin by generating 1D models, pathologically designed so that, at the last value of the discrete z vector, the integral of $\alpha_1(z)$ is approximately zero, which avoids having the end of the signal behave like a strong reflector. This is done with no loss of generality, since any data set can have such an addendum included beyond the deepest point of interest. The first such model has 4 wavespeeds, the reference $c_0 = 1500\text{m/s}$, and $c_1 = 1600\text{m/s}$, $c_2 = 1650\text{m/s}$, $c_3 = 1467\text{m/s}$. (This last wavespeed value ensures that the above constraint holds.) These latter 3 wavespeeds correspond to layers which begin at depths 300m, 500m, and 700m respectively.

This model is used to generate full-bandwidth data ($D(z)$ at pseudo-depth z), then the Born approximation $\alpha_1(z)$, and its integral, which we refer to as $\mathcal{H}\{\alpha_1\}$. These are plotted respectively in Figure 1. This same process is carried out on a suite of Earth models, each having been chosen for (i) simplicity, and/or (ii) high-contrast, and/or (iii) somewhat complex structure. Table 1 details the models used.

1.2 Brute Implementation

We begin by “naive” application of the formula. The first term returns α_1 . The result of computing and adding-in the second and third terms is seen in Figure 2. The figure is organized as follows: the top panel (a) is the synthetic data; below this (b) consists of two functions, the Born approximation $\alpha_1(z)$ (dashed), and the true perturbation $\alpha(z)$ (dotted). The required tasks of the inverse series are clear: the inversion must correct the amplitudes, and the imaging must correct the locations. In other words, the inversion must make the dashed be the same as the dotted in the *up-down* direction of the plots, and the imaging must




Figure 1: Synthetic data (a) corresponding to Model 1 in Table 1 and its integrals; the Born approximation $\alpha_1(z)$ is in (b), and its integral $H\{\alpha_1\}$ is in (c). This last is the main ingredient in computing the coupled imaging and inversion. All three plots are against pseudo-depth z (m).

Depth (m)	Model 1 (m/s)	Model 2 (m/s)	Model 3 (m/s)	Model 4 (m/s)
300-500	1600	1600	2000	2000
500-700	1650	1650	1700	2200
700-750	1467	1600	1422	1423
750-800	–	1570	–	–
800-870	–	1530	–	–
870-910	–	1500	–	–
910- ∞	–	1454	–	–

Table 1: All Earth models used in the following imaging/inversion examples. All have reference media ($z < 300$ m) characterized by wavespeed $c_0 = 1500$ m/s. Model 2 has structure deeper than the others: the dash – signifies that the model remains constant at the last given wavespeed value. For instance, Model 1 is constant at 1467m/s below 700m.

make the dashed be the same as the dotted in the *left-right* direction. The next panel (c) again illustrates $\alpha_1(z)$ (dashed), and includes the cumulative result of the terms in equation (3) beyond the first (solid). The lower panel (d) superimposes the full inversion results, $\alpha_1 +$ cumulative result (solid), against the true perturbation $\alpha(z)$. In all figures of this kind that follow, “added value” associated with higher order terms in the series is demonstrated by having the plots in (d) become close to one another.

One can see the disturbances created by the series at the discontinuities – clearly more terms are needed to correct the location of the discontinuities. It is interesting to note that, by the third term, the amplitude correction (inversion) has come, visually, close to accomplishing its task; away from the discontinuities, and following the Born structure of the model, the layers have found their desired amplitude.




Figure 2: The third-order correction from equation (3), using Model 1 from Table 1. (a) Data input; (b) Born approximation (dashed) vs. true perturbation (dotted); (c) Born approximation (dashed) vs. second-order correction (solid); (d) sum of Born approximation and correction (solid) vs. true perturbation. The inversion task is close to being done.

Continuing with the naive application of equation (3), compute the fourth term and add.

This is plotted in Figure 3. Clearly the discontinuities on the higher derivative operators quickly create large oscillations over the whole signal. Shortly hereafter the sum “blows up”. Something more sophisticated is required.




Figure 3: The cumulative sum up to fourth order from equation (3), using Model 1 from Table 1. (a) Data input; (b) Born approximation (dashed) vs. true perturbation (dotted); (c) Born approximation (dashed) vs. second-order correction (solid); (d) sum of Born approximation and correction (solid) vs. true perturbation. Divergence occurs shortly hereafter.

1.3 Stabilizing the n 'th Derivative

The uncontrolled oscillation of the unstable cumulative sum suggests that it is the high-frequency portions of the derivative operators which cause the problem. A simplistic regularization, amounting to a smoothed cutting-off of the highest frequencies, allows resolution to be traded-off for stability. The regularization works as follows. Windows in the k domain are constructed by convolving a gate function with a Gaussian; such a window is therefore defined with two parameters, the variance of the Gaussian, and the width of the gate.

The latter controls the frequency cut-off, and the former controls the smoothness of the “shoulders” of the cutoff. Figure 4 illustrates these windows.




Figure 4: *Four plots of regularization windows in the k domain for various parameters. A low variance Gaussian convolved with narrow gate is illustrated (a) vs. wide gate (b) as is a high variance Gaussian convolved with narrow gate (c) vs. wide gate (d). Character of the ensuing regularization may be altered by changing these parameters to suit specific situations.*




Figure 5: Regularized derivative operators in the k domain (solid) vs. un-regularized operators (dotted). (a) d/dz , (b) d^2/dz^2 , (c) d^3/dz^3 , and (d) d^4/dz^4 . “Levels” of regularization are dictated by the window parameters (see Figure 4).

In these numerical examples the derivatives are computed in the wavenumber domain. Figure 5 illustrates the un-regularized (dotted) and regularized derivative operators (solid) for orders 1 through 4.

Since the n 'th derivative is a linear, space invariant operation, its Fourier domain representation is equivalent to a singular value expansion. A weighted suppression of the Fourier coefficients of the operator at high frequency is therefore equivalent to stabilization via *truncated singular value decomposition*, in the parlance of linear inverse theory.

1.4 Numerical Examples of Simultaneous Imaging and Inversion

Using a sharp truncation (i.e. low variance) and a window that cuts off the high-frequency portion of the derivative operators on either end of the spectrum, we proceed to compute ≈ 100 terms in the series of equation (3) for the same input. The results are in Figure 6. The regularization parameters must be chosen for each example; we did this by trial and error for these examples.

Clearly much of the character of the true model is captured here – see the additional examples in Figures 7 – 9. The main deviation is in the large contrast examples, in which inaccuracy




Figure 6: The cumulative sum of ≈ 100 terms from equation (3), using Model 1 from Table 1. (a) Data input; (b) Born approximation (dashed) vs. true perturbation (dotted); (c) Born approximation (dashed) vs. (solid); (d) sum of Born approximation and correction (solid) vs. true perturbation. Using a low-variance Gaussian and a gate that cuts the derivative operator off on each end of the spectrum, 100 terms from equation (3) are computed and summed. Results capture closely the desired result.

in the imaging (reflector location) appears. In these examples, as evidenced by the poorer resolution, the derivative operators must also be more aggressively regularized (truncated); however, the missing bandwidth does not explain the inaccuracy. It is reasonable to postulate that it is due to the missing higher order imaging terms in the approximation; indeed there was earlier evidence that the partial nature of this scheme meant inaccuracy at higher reflection coefficients.




Figure 7: The cumulative effects of ≈ 100 terms from equation (3), using Model 2 from Table 1. (a) Data input; (b) Born approximation (dashed) vs. true perturbation (dotted); (c) Born approximation (dashed) vs. correction (solid); (d) sum of Born approximation and correction (solid) vs. true perturbation. Using a low-variance Gaussian and a gate that cuts the derivative operator off on each end of the spectrum, 100 terms from equation (3) are computed and summed. Results capture closely the desired result.




Figure 8: The cumulative effects of ≈ 100 terms from equation (3), using Model 3 from Table 1. (a) Data input; (b) Born approximation (dashed) vs. true perturbation (dotted); (c) Born approximation (dashed) vs. correction (solid); (d) sum of Born approximation and correction (solid) vs. true perturbation. Using a low-variance Gaussian and a gate that cuts the derivative operator off on each end of the spectrum, 100 terms from equation (3) are computed and summed. Results capture closely the desired result, but some inaccuracy in the high-contrast correction is noticed.




Figure 9: The cumulative effects of ≈ 100 terms from equation (3), using Model 4 from Table 1. (a) Data input; (b) Born approximation (dashed) vs. true perturbation (dotted); (c) Born approximation (dashed) vs. correction (solid); (d) sum of Born approximation and correction (solid) vs. true perturbation. Using a low-variance Gaussian and a gate that cuts the derivative operator off on each end of the spectrum, 100 terms from equation (3) are computed and summed. Results capture closely the desired result, but greater inaccuracy in the high-contrast correction is noticed.

2 Noise and Bandlimitation

In this section we consider two key departures of real-world problems of seismic processing from theory: noise and bandlimitation.

2.1 Robustness to Incoherent Noise

It has been mentioned that, in principle, equation (3) should be unstable, and, indeed, it has been found that the derivative operators require high-frequency truncation, more-so with greater model contrast. Once this has been accomplished the recovered models capture the sharpness of the contrasts admirably, through the computation of ≈ 100 terms. These results apply for full bandwidth data with no noise. In the next section we will address the problem of bandlimitation. Here we consider the results in the presence of varying levels of additive incoherent (Gaussian) noise.

Consider the example of Figure 7. Three realizations of Gaussian noise (with variance of approximately %1 of the data amplitudes) are added to the data, and the imaging/inversion terms are recomputed using the $\int_0^z \alpha_1(z') dz'$ that comes from integrating the noisy data. Figures 10 – 12 contain the results.

The results are clearly deteriorated by the presence of noise – in fact, *realization to realization* the same level of noise can produce quite different results. Both the location and the amplitude of the correction are corrupted, but a qualitative conclusion is that the inversion, or amplitude results are the most sensitive. In all cases the imaging component, i.e. the movement of the reflectors, still marks a great improvement over the Born approximation.

The quality of these Earth models, recovered in the presence of noise, is heartening in the sense that one might expect the smallest amount of noise to render computation of equation (3) completely unstable. Nevertheless, it is clear that a very high-fidelity estimate of α_1 will be of tremendous value. The estimation of such an input might include edge-preserving noise reduction strategies, as well as (presumably) low frequency/wavenumber filling strategies.




Figure 10: The data used to generate Figure 7 is corrupted with %1 noise, and the imaging/inversion results are recomputed. Organization of results is as in Figure 7. This is the first of three realizations.




Figure 11: The data used to generate Figure 7 is corrupted with %1 noise, and the imaging/inversion results are recomputed. Organization of results is as in Figure 7. This is the second of three realizations.




Figure 12: The data used to generate Figure 7 is corrupted with %1 noise, and the imaging/inversion results are recomputed. Organization of results is as in Figure 7. This is the third of three realizations.

2.2 Spectral Extrapolation: Computing α_1 from Bandlimited Data

There are some striking computational similarities between these inverse scattering subseries methods and other well-studied problems of inversion of seismic wave field measurements. In particular, at its core are integral operations on a quantity which ideally should be realized over a full band of frequencies, including a DC component. Since this quantity is based on wave field measurements which are unavoidably bandlimited, however, it too is bandlimited in practice. The consequences to the output of the subseries requires study.

As with seismic inverse methods such as impedance inversion, methods for extension of the spectrum to zero frequency do exist. In this section an investigation of the applicability of such methods to our various subseries, in their current 1D incarnations, is presented. We show some specific results and their application to examples akin to those previously presented; but the idea here is to broadly propose the use of a framework, long in the literature (see below for references), for tackling the important problem of bandlimited inversion.

A measured signal, wavelet deconvolved, may be assumed to be the bandlimited expression of some reasonably simple full-bandwidth signal type. An example might be a series of lagged and scaled delta functions, i.e. a reflectivity series. The spectrum of the measured signal – which is the full bandwidth spectrum multiplied by a gate function – can be extended, in principle uniquely, guided by this assumption.

The calculation of α_1 is linear inversion, and with the right assumptions it is equivalent to 1D inversion for acoustic impedance. So the root of the low-frequency trouble is in essence the same for the inverse series as it is for impedance inversion; strategies for coping with this lack in impedance recovery may therefore be readily applicable in the 1D normal incidence inverse series examples presented here.

Basic Concepts of Bandlimited Impedance Inversion

We begin by discussing spectral extrapolation proper, as part of a brief review of methods for bandlimited inversion. A delta function in the time domain has a complex sinusoid as a spectrum. The frequency of this sinusoid is uniquely determined by the lag of the delta function. Since a sinusoid with additive noise is adequately described with autoregressive (and correctly modelled with ARMA) models, a spectrum of this kind that has been bandlimited can be effectively predicted and therefore extended beyond the bandwidth. Larger order ARMA and AR models may, further, predict the summation of many sinusoids (of different frequencies), and therefore extend the spectra of bandlimited delta series with many varied lags. Consider the reflectivity series that is this sum of delta functions, following Walker and Ulrych (1983):

$$r(t) = \sum_k b_k \delta(t - k\Delta), \quad (4)$$

where b_k are the reflectivity coefficients and Δ is the time interval of the experiment. The Fourier transform, or the spectrum of the reflectivity, $R(f)$, is:

$$\begin{aligned} R(f) &= \int_{-\infty}^{\infty} e^{-i2\pi ft} \left[\sum_k b_k \delta(t - k\Delta) \right] dt \\ &= \sum_k b_k e^{-i2\pi fk}, \end{aligned} \quad (5)$$

if Δ is set to unity. So the $R(f)$ is this sum of weighted complex sinusoids. With a finite bandwidth, the data $R(f)$ may be extended to the lower (and, to some extent, the higher) frequencies via a prediction scheme. Although the noisy sum of sinusoids is correctly modelled via an ARMA model, Walker and Ulrych (1983) recommend, rather, a truncated AR prediction approach as being more stable.

An autoregressive process has a predictable part and an unpredictable, or innovational, part. That is, the j 'th element of a series y_j can be written as a linear combination of previous elements, plus an error (innovation) term:

$$y_j = \sum_{k=1}^p a_k y_{j-k} + e_j, \quad (6)$$

or

$$y_j = - \sum_{k=1}^p g_k y_{j-k} + e_j, \quad (7)$$

where $g_k = -a_k$ and $g_0 = 1$ are the coefficients of the prediction error filter. With knowledge of the variance of the error e_j (σ_e), and having in some appropriate way computed the autocorrelation matrix of y_j (\mathbf{R}_{yy}), the prediction error filter \mathbf{g} is found by solving the system

$$\mathbf{R}_{yy}\mathbf{g} = \sigma_e^2 \mathbf{i}, \quad (8)$$

where $\mathbf{i} = [1, 0, 0, \dots, 0]^T$ and \mathbf{g} and \mathbf{R}_{yy} etc. all have complex elements.

Associating bandlimited $R(f)$ with y_j , the filter \mathbf{g} may be applied to existing values of $R(f)$ to produce an extended spectrum data point, one discrete frequency step closer to zero. This point is thereafter treated as an existing value of the reflectivity spectrum, and the process is repeated.

Spectral extension, which is essentially extrapolation via a difference equation, has a growing error in practical application as it predicts further from the known data values. There is in fact an overall tendency of extended spectra to decay. This can be avoided by an efficient

alternate form of prediction based on the gap-filling algorithms of Wiggins (1972), in which the existence of known data values on *both* sides of the unknown region is utilized to avoid such accumulating error. Since, by conjugate symmetry, the negative frequencies are also known, the missing lower frequencies may be thought of as such a gap, on either side of which are known data. An unknown point in the spectrum is considered to be the weighted sum of predictions from the “left” and from the “right” data sets (and in which the weights are determined by minimizing the overall prediction error). Walker and Ulrych (1983) and Ulrych and Walker (1984) show that a scheme of this kind finds all the missing data points (of the discrete spectrum) at once. One may extend to high frequencies in the same way, although, since the high frequency gap is considerably larger than the low frequency one, the predictions will be more highly attenuated.

The problem may, secondly, be addressed by utilizing the formalism of inverse theory (Oldenburg, 1984; Oldenburg et al., 1983); the approach adopted here is the work of Ulrych (1989), which uses an entropic norm to solve for the whole spectrum with the existing spectrum as the input. The bandlimited signal is viewed as being non-unique, in the sense that an infinite number of *models* (i.e. time-domain signals) may be Fourier transformed and bandlimited to produce the *data* (i.e. the bandlimited Fourier transform). The method incorporates prior information garnered from the data (usually a threshold) as a constraint, which directs the model “choice”, from the set of allowable models, to be that which is sparse, or spike-like. In such models the low and high frequencies tend to be present.

Either of the two approaches discussed above is well-suited to the task of estimating the expected form of the full-bandwidth Born approximation. In the examples to follow, the latter is used, simply because it tends to estimate the high frequencies of the spectrum with greater accuracy.

Numerical Examples Using Bandlimited Data

In Figure 13, detail of the filled data is given close to a single reflection. The top panel (a) is the full bandwidth data, the second panel (b) is the bandlimited equivalent (assuming $\Delta t = 0.004\text{s}$, these examples correspond to a band of 10–50Hz), the third is the spectrally-extrapolated output, and the fourth is a detail (zoom) of the extrapolated (solid) pulse vs. the full bandwidth pulse (dashed).

In Figure 14, the results of Figure 7 are re-computed using spectrally-extrapolated estimates of the data, with only bandlimited traces as input. Further panels are added to the top of this figure; first is the bandlimited data, below that is the filled data, and below that are the panels illustrating the inversion/imaging as previously used.

In Figures 14 and 15, the relatively spatially-complex data example and the high-contrast data example, both bandlimited are shown to respond very well to the spectral-extrapolation pre-processing.

Of course, these examples have not specifically pushed the limits of spectral extrapolation technology. The key is to, in some quantitative way, address the issue of bandlimitation

in these non-linear inversion schemes. The value of these examples is in the demonstration of working methods for compensating, often via some reasonable prior knowledge, for imperfections of the seismic experiment.




Figure 13: Detail of the filled data is given close to a single reflection. (a) full bandwidth data, (b) the bandlimited equivalent, (c) spectrally-extrapolated output, and (d) detail of the extrapolated (dashed) pulse vs. the full bandwidth pulse (solid).




Figure 14: *Simultaneous imaging and inversion as applied to Model 2 from Table 1 is repeated on spectrally-extrapolated estimates of the data, using only bandlimited traces as input. Further panels are added to the top of this figure (a) the input bandlimited data, (b) the spectrally-extrapolated (recovered) full-bandwidth data, (c) Born approximation (dashed) vs. true perturbation (dotted); (d) Born approximation (dashed) vs. second-order correction (solid); (e) sum of Born approximation and correction (solid) vs. true perturbation.*




Figure 15: *Simultaneous imaging and inversion as applied to Model 4 (high contrast) from Table 1 is repeated on spectrally-extrapolated estimates of the data, using only bandlimited traces as input. Further panels are added to the top of this figure (a) the input bandlimited data, (b) the spectrally-extrapolated (recovered) full-bandwidth data, (c) Born approximation (dashed) vs. true perturbation (dotted); (d) Born approximation (dashed) vs. second-order correction (solid); (e) sum of Born approximation and correction (solid) vs. true perturbation.*

3 The Relationship Between *Order* and *Frequency Content*

In low-order truncations of the simultaneous imaging and inversion subseries there occurs an “explosion” in the depth domain – uncontrolled high-amplitude oscillation of the reconstruction (see for example Figure 3). After the addition of many terms, this oscillation can be seen to “settle down” upon the desired reconstruction. (This characteristic evolution-with-order is also true of the leading order imaging subseries.) It is tempting to ignore this behaviour prior to convergence, since it appears to be uninterpretable. Here we investigate the numerical convergence issues of the simultaneous imaging and inversion subseries using some simple Taylor’s series examples as a guide. The result is an analytical/numerical framework within which to better understand the seemingly unstable order-by-order behaviour of subseries that locate reflectors.

We know that the n ’th term in the simultaneous imaging and inversion subseries (and the leading order imaging subseries) involves the n ’th spatial derivative of an input. We may therefore make the very loose mathematical comment that, in the wavenumber domain, terms in the series will involve increasing powers of (ik) :

$$\alpha_{SII}(k) \approx \dots + (ikC)^n G(k) + \dots \quad (9)$$

and also some function $G(k)$, and some constant C . Since we have *chosen* the true model in these tests, we also know something about what we are constructing; namely, discontinuous functions that correct spatial locations of reflectors and their amplitudes (in this case the discontinuous functions are Heaviside functions). In the wavenumber domain, then, the reconstruction will be an exponential function with a wavenumber-dependent weight, so again loosely:

$$\alpha_{SII}(k) \approx F(k)e^{ikC} \quad (10)$$

where $F(k)$ is the weighting function.

In other words, the simultaneous imaging and inversion process is very similar to the Taylor’s series expansion of an exponential function about zero, i.e. involving an infinite series of polynomials:

$$e^x = 1 - x + \frac{1}{2}x^2 - \frac{1}{3!}x^3 + \dots, \quad (11)$$

where we interpret the argument x in our case as being ikC .

Since our imaging/inversion goal is to construct these corrective, discontinuous signals numerically over a finite wavenumber interval, the correct analogy is that of constructing e^x over some fixed interval $x = (0, x_{max})$. In Figure 16, we estimate the (better behaved) function e^{-x} over such a fixed interval for a number of orders.

The evolution of the approximation with truncation order is straightforward: for each successive term added, the approximation is made accurate (i.e. it converges) over a larger interval of x . Beyond this region of convergence the approximation diverges at a rate of x^n for a truncation at the n 'th term.




Figure 16: Taylor's series approximation of e^{-x} over a fixed interval; six approximations, 0'th-3'rd, followed by 9'th and 17'th. As more terms are added, the approximations converge over a lengthening interval of x .

Applying this thinking to the simultaneous imaging and inversion subseries (and by association the general nature of the inverse scattering series terms for processing and inversion of primaries), we recognize that there is an implied relationship between the convergent bandwidth of the reconstruction and the order at which the series construction is truncated. To wit: since the interval of construction (of exponential-like models) is in the k domain, we can expect the series to converge, with added terms, over a growing wavenumber interval. To demonstrate, we include in Figures 17 – 21 numeric examples of the reconstruction of α_{SII} in the k domain (Figure 17a etc.). Although the signal in the k domain is somewhat more complicated than e^{-x} in the x domain, the evolution across increasing truncation orders is seen to behave very similarly, with an increasing interval of convergence in k . As a final

step for each example (i.e. *after* each summation is complete), we manually suppress the divergent frequencies from each example, inverse Fourier transform, and display the result (Figures 17b, c – 21b, c). The reconstructions go from smooth to sharp as we move from low-order to high-order truncation.

The direct relationship between the frequency content of the reconstruction is in this way made more apparent to the eye (more apparent, that is, than if we had left in the dominant but uninformative divergent portions of the spectra) and the approximation order of the model. This insight may be of use as it implies a possible tradeoff that can be made between the resolution of the reconstruction and the number of terms used (which in more realistic instances could have very important computational consequences). It has not escaped the authors' attention that the very bandwidth extrapolation methods discussed in this paper (to deal with missing low frequency data) may be applicable in such series computations to reduce the number of terms needed: i.e., the series terms might be used up to a certain order, beyond which the remaining signal on the k interval may be “predictable” in the sense of bandwidth extrapolation. This remains speculative.




Figure 17: Truncated computation of $\alpha_{SII}(z)$, order 5: (a) in the wavenumber domain, the construction appears to diverge rapidly with larger k ; (b) and (c) manually suppressing the divergent wavenumbers and inverse Fourier transforming, we see that the correction is present but highly smoothed when compared to the true perturbation (dotted) and the input Born approximation (dashed).




Figure 18: Truncated computation of $\alpha_{SII}(z)$, order 10: (a) in the wavenumber domain, the construction appears to diverge rapidly with larger k , although less so at lower k than for order 5; (b) and (c) manually suppressing the divergent wavenumbers and inverse Fourier transforming, we see that the correction is somewhat less smoothed when compared to the true perturbation (dotted) and the input Born approximation (dashed).




Figure 19: Truncated computation of $\alpha_{SII}(z)$, order 20: (a) in the wavenumber domain, the construction can now be seen to be non-increasing over a wider interval of k than in the 5'th and 10'th order truncations; (b) and (c) manually suppressing the divergent wavenumbers and inverse Fourier transforming, we see that the correction is concurrently becoming “sharper” when compared to the true perturbation (dotted) and the input Born approximation (dashed).




Figure 20: Truncated computation of $\alpha_{SII}(z)$, order 40: (a) in the wavenumber domain, the construction can now again be seen to be non-increasing over a wider interval; (b) and (c) manually suppressing the divergent wavenumbers and inverse Fourier transforming, we see again that the correction is concurrently becoming “sharper” when compared to the true perturbation (dotted) and the input Born approximation (dashed).




Figure 21: Truncated computation of $\alpha_{SII}(z)$, order 80: (a) in the wavenumber domain, the construction can now again be seen to be non-increasing over a wider interval; (b) and (c) manually suppressing the divergent wavenumbers and inverse Fourier transforming, we see again that the correction is concurrently becoming “sharper” when compared to the true perturbation (dotted) and the input Born approximation (dashed).

4 Conclusions

We consider some basic numeric issues associated with the computation and interpretation of a subseries that is involved with both imaging and inversion of seismic primaries. The subseries and the computations are concerned with the simple 1D normal incidence acoustic problem.

The raw computation of truncated versions of the subseries are shown to produce unstable results, a problem that may be overcome through the use of approximate derivative operators whose highest frequency components are suppressed. The operator approximations are designed to accommodate differing contrasts in the data and – we show – incoherent noise.

We also demonstrate the use of a form of spectral extrapolation that exists in the literature ostensibly for bandlimited impedance inversion. This methodology, on the assumption of delta-like data events, extrapolates from existing bandwidth intervals to missing intervals. We show with some simple examples the application of this approach to the computation of the Born approximation – filling in the low frequencies of the linear result provides a full-bandwidth input to the higher order terms in the series.

Finally, we consider the numerics of *truncated* model estimates from the imaging-inversion subseries; we demonstrate the relationship between truncation order and frequency/wavenumber content of the constructed model using a simple Taylor’s series expansion of an exponential function as a guide.

The purpose of this work is to sketch out some of the basic numerical issues at the core of the inverse scattering subseries’ for the imaging and inversion of primaries. The expectation is that some or all of the issues encountered here will be visited on inverse scattering series-based computations in higher dimensions and for more complicated models; further, that some of the strategies developed for these simple instances might be extended (also to higher dimensions and greater complexity).

Acknowledgments

The authors thank the sponsors and members of M-OSRP and CDSST, especially Simon Shaw, and Bogdan Nita, for discussion and support.

References

Innanen, K. A. “Methods for the Treatment of Acoustic and Absorptive/Dispersive Wave Field Measurements.” *Ph.D. Thesis, University of British Columbia*, (2003).

Innanen, K. A. and Weglein, A. B., “Simultaneous Imaging and Inversion with the Inverse Scattering Series.” *Proceedings of the Eighth International Congress of the SBGf and Fifth Latin American Geophysical Conference, Rio de Janeiro, Brazil*, (2003).

Innanen, K. A., Nita, B., and Weglein, A. B. “Investigating the Grouping of Inverse Scattering Series Terms: Simultaneous Imaging and Inversion I (Theory).” *M-OSRP Annual Report*, (2004): This Report.

Oldenburg, D. W., Scheuer, T., and Levy, S., “Recovery of the Acoustic Impedance from Reflection Seismograms.” *Geophysics* 48, (1983): 1318–1337.

Oldenburg, D. W., “An Introduction to Linear Inverse Theory.” *IEEE Transactions on Geoscience and Remote Sensing GE-22*, (1984): 665–674.

Shaw, S. A., and Weglein, A. B. “A leading order imaging series for prestack data acquired over a laterally invariant acoustic medium. Part II: Analysis for data missing low frequencies.” *M-OSRP Annual Report*, (2004): This Report.

Ulrych, T. J. “Minimum Relative Entropy and Inversion.” *In: Geophysical Inversion, SIAM, J. B. Bednar, L. R. Lines, R. H. Stolt, A. B. Weglein*, (1989): 158.

Ulrych, T. J., and Walker, C., “On a Modified Algorithm for the Autoregressive Recovery of the Acoustic Impedance.” *Geophysics* 48, (1984): 2190–2192.

Walker, C., and Ulrych, T. J., “Autoregressive Recovery of the Acoustic Impedance.” *Geophysics* 48, (1983): 1338–1350.

Weglein, A. B., Araujo, F. A., Carvalho, P. M., Stolt, R. H., Matson, K. H., Coates, R., Foster, D. J., Shaw, S. A., and Zhang, H. “Topical Review: Inverse-scattering Series and Seismic Exploration.” *Inverse Problems* 19, (2003):R27–R83.

Wiggins, R. A., and Miller, S. P., “New Noise Reduction Techniques Applied to Long Period Oscillations of the Alaskan Earthquake.” *BSSA* 62, (1972): 417–479.