

Geoacoustic inversion by mode amplitude perturbation

Travis L. Poole^{a)}

Woods Hole Oceanographic Institution, Gigelow 209, MS#12, Woods Hole, Massachusetts 02543

George V. Frisk^{b)}

Department of Ocean Engineering, Florida Atlantic University, SeaTech Campus, 101 North Beach Road, Dania Beach, Florida 33004

James F. Lynch^{c)}

Woods Hole Oceanographic Institution, Bigelow 209, MS#12, Woods Hole, Massachusetts 02543

Allan D. Pierce^{d)}

Aerospace and Mechanical Engineering, Boston University, Boston, Massachusetts 02215

(Received 13 August 2007; accepted 13 November 2007)

This paper introduces a perturbative inversion algorithm for determining sea floor acoustic properties, which uses modal amplitudes as input data. Perturbative inverse methods have been used in the past to estimate bottom acoustic properties in sediments, but up to this point these methods have used only the modal eigenvalues as input data. As with previous perturbative inversion methods, the one developed in this paper solves the nonlinear inverse problem using a series of approximate, linear steps. Examples of the method applied to synthetic and experimental data are provided to demonstrate the method's feasibility. Finally, it is shown that modal eigenvalue *and* amplitude perturbation can be combined into a single inversion algorithm that uses all of the potentially available modal data. © 2008 Acoustical Society of America. [DOI: 10.1121/1.2821975]

PACS number(s): 43.30.Bp, 43.30.Ma [AIT]

Pages: 667–678

I. INTRODUCTION

In 1987 Rajan *et al.*¹ described a perturbative inversion method for obtaining the geoacoustic properties of the seabed from measurements of modal eigenvalues. Since the bottom influences the modal amplitudes as well, it seems natural to ask if the modal amplitudes can also be used to estimate the bottom properties. Furthermore, perhaps both the modal eigenvalue and amplitude data can be used together in a single inversion algorithm to achieve an improved result. This paper addresses both of these issues. Its main objective is to introduce a perturbative inversion scheme analogous to that in Ref. 1 using modal amplitudes, instead of modal eigenvalues, as the input data. We then show that modal eigenvalue *and* amplitude perturbation can be combined into a single algorithm, thereby combining the advantages of both methods.

The method assumes that measurements of the acoustic pressure field due to a cw point source have been made at one or more low frequencies (<500 Hz) in the water column. The quantity of primary interest is the compressional wave speed as a function of depth, but the density and attenuation profiles can also be sought, and we derive the nec-

essary equations for determining them as well. The effects of shear waves are ignored in this development, but they could be included in a more general formalism.

Unlike Rajan *et al.*,¹ who made explicit use of the Hankel transform method for estimating the modal eigenvalues, we intentionally avoid a detailed discussion of the manner in which the data are collected and processed to obtain the required modal information. Though this is a critical step along the path to determining the bottom parameters, it is a separate issue from the modal amplitude inversion itself. We therefore assume knowledge of the modal amplitudes (preferably with an estimate of the uncertainty in those data) as well as information about the water column sound speed profiles at the locations where the bottom properties are to be determined. Whether these estimates are the output of a Hankel transform of horizontal line array (HLA) data collected over range as in Ref. 1, or from filtering the modes on a vertical line array (VLA) over depth as suggested by Ref. 2, or by some other method (for example, Ref. 3), the methodology of the modal amplitude inversion will remain the same. Thus, we avoid in-depth discussion of how the data are collected, and focus instead on the subsequent use of the data to estimate the bottom properties. Our preferred method of obtaining the modal amplitude data from measurements of the point source acoustic field will be the subject of a future publication.

Perturbative inversion methods approach the solution of the nonlinear problem of determining the bottom parameters from measurements of the acoustic field via a series of small, linear steps. During the first step of the algorithm a user-provided background model is used to calculate the parameters of the acoustic field (such as the eigenvalues or modal

^{a)}Present address: Grant Institute of Earth Science, University of Edinburgh, Kings Buildings, Edinburgh, EH9 3JW, United Kingdom. Electronic mail: travis.poole@ed.ac.uk

^{b)}Electronic mail: gfrisk@seatech.fau.edu

^{c)}Electronic mail: jlynch@whoi.edu

^{d)}Electronic mail: adp@bu.edu

amplitudes) using a forward model. Using this background model, the first-order derivatives of the acoustic parameters with respect to the bottom parameters are also computed. Next, the computed acoustic parameters are compared to the measured values, and we compute the change to the background model needed to correct (in a least-squares sense) our estimates based on an assumption of linearity. Since this linear assumption is not strictly correct, our correction will not give us the true bottom properties. If the background model is sufficiently similar to the true bottom, however, the correction will give us an improved estimate of the geoaoustic properties which can be used as a new background model, and the process can be repeated. In theory, each iteration should bring us closer to the true set of bottom parameters, and we can accept the result to which the process converges to be the best estimate of those parameters. This iterative sequence of steps, which typically involves at most a few tens of iterations, is to be contrasted with matched field inversion methods which search the parameter space by computing thousands or tens of thousands of forward models.

The paper is organized as follows: In Sec. II we derive the perturbation results we use in the method and set up the linear equations to be inverted. Section III discusses the singular value decomposition method of performing the inversion, and discusses the errors associated with such an inversion. In Sec. IV we apply the mode amplitude perturbation algorithm to two data sets to demonstrate its performance. The first is a simple synthetic case in which the true answer is known, and the second is a real data set from the LWAD 99-1 Experiment.⁴ In Sec. V we combine the eigenvalue and mode amplitude perturbation methods into a single algorithm and apply this combined method to a synthetic data set to investigate its capabilities. The last section summarizes our results and conclusions.

II. DERIVATION OF THE MODAL AMPLITUDE PERTURBATIVE INVERSION ALGORITHM

In this section we derive the modal amplitude perturbative inversion algorithm, starting with the application of perturbation theory to the basic normal mode equations. We then discuss various approaches to bottom parametrization, and end the section by describing the inversion process itself. Our derivation of the perturbation result for the mode amplitudes follows the work of Tindle *et al.*⁵

We start with the depth-separated normal mode equation (cf. Refs. 6–8)

$$\frac{1}{\rho(z)} \frac{d^2 Z_n}{dz^2} + \frac{d}{dz} \left(\frac{1}{\rho(z)} \right) \frac{dZ_n}{dz} + \frac{1}{\rho(z)} [q(z) - k_n^2] Z_n = 0, \quad (1)$$

where z is the depth coordinate, positive downwards (with the air-water interface at $z=0$), $\rho(z)$ is the density, $Z_n(z)$ is the n th mode function, k_n is the modal eigenvalue associated with $Z_n(z)$, and $q(z) = k^2(z) = \omega^2 / [c^2(z)]$, where $c(z)$ is the sound speed, and ω is the angular frequency of the time-harmonic source signal. We propose a perturbation to the waveguide such that $q(z) \rightarrow q(z) + \Delta q(z)$, which causes perturbations in the other terms: $Z_n(z) \rightarrow Z_n(z) + \Delta Z_n(z)$, and $k_n \rightarrow k_n + \Delta k_n$.

If we collect the unperturbed terms we get the unperturbed equation. If we collect terms with first-order perturbations we find that

$$\begin{aligned} \frac{1}{\rho(z)} \frac{d^2 \Delta Z_n}{dz^2} + \frac{d}{dz} \left(\frac{1}{\rho(z)} \right) \frac{d \Delta Z_n}{dz} + \frac{1}{\rho(z)} [q(z) - k_n^2] \Delta Z_n \\ + \frac{1}{\rho(z)} [\Delta q(z) - 2k_n \Delta k_n] Z_n = 0. \end{aligned} \quad (2)$$

We assume that the perturbations are small, so terms of higher than first order are neglected.

Because the unperturbed normal modes form a complete set, we can expand any function of z in terms of them. We propose an expansion of ΔZ_n of the form $\Delta Z_n(z) = \sum_j a_{nj} Z_j(z)$. Substituting this into Eq. (2) we find that

$$\begin{aligned} \sum_j a_{nj} \left[\frac{1}{\rho} \frac{d^2 Z_j}{dz^2} + \frac{d}{dz} \left(\frac{1}{\rho} \right) \frac{dZ_j}{dz} \right] + \frac{1}{\rho} (q(z) - k_n^2) \sum_j a_{nj} Z_j \\ + \frac{1}{\rho} (\Delta q(z) - 2k_n \Delta k_n) Z_n = 0. \end{aligned} \quad (3)$$

The term in square brackets can be replaced using Eq. (1):

$$\begin{aligned} \sum_j a_{nj} \left[-\frac{1}{\rho} (q(z) - k_j^2) Z_j \right] + \frac{1}{\rho} [q(z) - k_n^2] \sum_j a_{nj} Z_j \\ + \frac{1}{\rho} [\Delta q(z) - 2k_n \Delta k_n] Z_n = 0. \end{aligned} \quad (4)$$

Some of the terms containing $q(z)$ cancel, leaving us with

$$\sum_j a_{nj} (k_j^2 - k_n^2) \frac{Z_j}{\rho} + \frac{1}{\rho} (\Delta q(z) - 2k_n \Delta k_n) Z_n = 0. \quad (5)$$

From this point we can make use of the orthonormality property of the normal modes. If we apply the operator $\int_0^D (\bullet) Z_n(z) dz$ to the equation, we are left with

$$\begin{aligned} \int_0^D \frac{\Delta q(z) Z_n^2(z)}{\rho(z)} dz - 2k_n \Delta k_n \\ = 0 \Rightarrow \Delta k_n = \frac{1}{2k_n} \int_0^D \frac{\Delta q(z) Z_n^2(z)}{\rho(z)} dz, \end{aligned} \quad (6)$$

which is the result in Ref. 1. If we instead apply the operator $\int_0^D (\bullet) Z_m(z) dz$ (note the change of subscript from n to m) we get

$$\begin{aligned} a_{nm} (k_m^2 - k_n^2) + \int_0^D \frac{\Delta q(z) Z_m(z) Z_n(z)}{\rho(z)} dz \\ = 0 \Rightarrow a_{nm} = \frac{1}{k_n^2 - k_m^2} \int_0^D \frac{\Delta q(z) Z_n(z) Z_m(z)}{\rho(z)} dz, \end{aligned} \quad (7)$$

which is valid for $m \neq n$. This result, along with the above-used expansion, allows us to express the change in a mode function $\Delta Z_n(z)$ due to a change in the profile, $\Delta q(z)$. Since the quantity that is actually changing in the bottom is the sound speed, $c(z)$, we can use the perturbation result that $\Delta q(z) = -2\Delta c(z) \omega^2 / c^3(z)$ to put the expression in terms of the sound speed perturbation. It should be noted that Ref. 5 lists

the $m=n$ term as being $a_{nn}=-1/2\sum_{m\neq n}a_{nm}^2$ based on the idea that the mode function must retain its normalization. However, this result is inconsistent with the earlier neglect of terms of second and higher order. It can be shown, however, that the actual value for the $n=m$ term must be zero, since the change of the mode function must be orthogonal to the mode function itself:

$$\int_0^D \frac{Z_n^2(z)}{\rho(z)} dz = 1 \Rightarrow \frac{\partial}{\partial X} \int_0^D \frac{Z_n^2(z)}{\rho(z)} dz = \frac{\partial}{\partial X} 1$$

$$\Rightarrow \int_0^D \frac{2Z_n(z) \partial Z_n(z)/\partial X}{\rho(z)} dz = 0 \quad (8)$$

for any parameter X that would cause a change in the mode function. In order to properly make use of the normalization as in Ref. 5, one would have to retain higher order terms in the original perturbation.

The sound speed, $c(z)$, is not the only bottom parameter that affects the mode functions and eigenvalues. The density profile $\rho(z)$ is also a factor. Analogous to the results for sound speed in Ref. 5, we can use perturbation theory to determine the effects of perturbations to a density parameter. If we define the function $\beta(z)=1/\rho(z)$ and denote differentiation with respect to z with primes, our modal equation becomes

$$\beta(z)Z_n'' + \beta'(z)Z_n' + \beta(z)[k^2(z) - k_n^2]Z_n = 0. \quad (9)$$

Just as we did for sound speed, we can introduce a perturbation, collect first-order terms, expand the perturbed mode functions in terms of the unperturbed mode functions, and apply the orthonormality property of the unperturbed modes. When we do so, we obtain two perturbation results for the density parameter:

$$\Delta k_n = \frac{1}{2k_n} \int_0^D Z_n(z)Z_n'(z) \left(\Delta\beta \frac{\beta'}{\beta} + \Delta\beta' \right) dz, \quad (10)$$

which is the eigenvalue perturbation result for density, and

$$a_{nm} = \frac{1}{k_n^2 - k_m^2} \int_0^D Z_m(z)Z_n'(z) \left(\Delta\beta \frac{\beta'}{\beta} + \Delta\beta' \right) dz, \quad (11)$$

which is the result for the modal expansion coefficients, analogous to Eq. (7). Note that because both the equation for the eigenvalue perturbation, and the expansion of the mode function perturbation involve z derivatives, discontinuities must be handled with particular care.

Another bottom parameter of interest is the attenuation profile. Frisk⁶ and other texts show that if we introduce an imaginary component to the sound speed profile by making $k(z)$ complex, the eigenvalue becomes complex as well. While the change in the mode function itself is usually negligible when attenuation is added, the complex portion of the eigenvalue creates the appearance that the entire mode function has been reduced. If we make the small change $k(z) \rightarrow k(z) + i\alpha(z)$, the modal eigenvalue will be changed as well: $k_n \rightarrow k_n + i\delta_n$, where the modal attenuation is given by

$$\delta_n = \frac{1}{k_n} \int_0^D \frac{\alpha(z)}{\rho(z)} k(z) Z_n^2(z) dz. \quad (12)$$

This leads to the apparent change in the mode function $Z_n \rightarrow Z_n e^{-\delta_n r}$, or $\Delta Z_n = -Z_n(1 - e^{-\delta_n r})$. For many purposes it is often best to look at the apparent change in the mode function over a range step due to the attenuation within that range step. In such a case, the apparent change in the mode function from the beginning of the step to the end will be $Z_n \rightarrow Z_n e^{-\delta_n \Delta r}$, where Δr is the length of the range step. If $\delta_n \Delta r \ll 1$ we can use the Taylor expansion for the exponential, and keep only the linear terms, giving $Z_n \rightarrow Z_n(1 - \delta_n \Delta r)$ or $\Delta Z_n \approx Z_n \delta_n \Delta r$. Because of its dependence on range, attenuation can seriously complicate our inversion. Fortunately, in cases where the measurements are taken at sufficiently short ranges from the source, attenuation can be neglected, or estimated by other means.⁹

Further bottom parameters, such as shear speeds and shear attenuations, may also be sought, in which case perturbation results for those parameters would also be needed. For the applications to low-frequency sediment acoustics of interest in this paper, the results for sound speed, density, and attenuation perturbations are considered sufficient. In fact, we will focus on the sound speed perturbation results, since they tend to have the strongest effect on the modal parameters.

The next step in the derivation of the inversion algorithm is to discretize our representation of the bottom. We do this so that we can solve for a finite number of unknowns rather than functions of the continuous variable z . We start by making the assumption that the unknown functions $c(z)$, $\rho(z)$, and $\alpha(z)$ can be written in terms of a weighted expansion of some known depth functions with unknown coefficients. For example, $c(z)$ can be expanded as $c(z) = c_0(z) + \sum_i X_i c_i(z)$. Here $c_0(z)$ is a hypothesized, or background, model for the sound speed profile. The functions $c_i(z)$ are arbitrary, user-defined functions that should be selected so as to be able to capture the important features of the sound speed profile. The unknown scalar coefficients X_i are what we seek, since once we have them we can reconstruct the sound speed profile.

The number and complexity of the $c_i(z)$ functions that should be used depends on how detailed a profile is required, and what kinds of bottom features are considered possible. Ideally one should use the smallest number of parameters that fully capture the bottom features. However, since one does not usually know beforehand what bottom features are present, it will often be necessary to use more than the ideal number of parameters.

Once the unknown functions have been parametrized we can compute the derivatives of the mode functions with respect to the unknown scalars X_i using the perturbation results we derived earlier. To do this we make the substitution $\Delta c(z) = X_i c_i(z)$ for the sound speed equations, $\Delta\beta(z) = B_i \beta_i(z)$ for the density, and $\alpha(z) = A_i \alpha_i(z)$ for the attenuation, and use the fact that $\Delta Z_n / X_i \approx \partial Z_n / \partial X_i$ for sufficiently small values of X_i . The results of this procedure are

$$\frac{\partial Z_n(z)}{\partial X_i} = \sum_m a_{nm}^{X_i} Z_m(z), \quad (13)$$

where

$$a_{nm}^{X_i} = \frac{2\omega^2}{k_m^2 - k_n^2} \int_0^D \frac{c_i(z) Z_n(z) Z_m(z)}{c_0^3(z) \rho(z)} dz, \quad (14)$$

$$\frac{\partial k_n}{\partial X_i} = \frac{-1}{k_n} \int_0^D \frac{c_i(z) Z_n^2(z)}{c_0^3(z) \rho(z)} dz, \quad (15)$$

$$\frac{\partial Z_n(z)}{\partial B_i} = \sum_m a_{nm}^{B_i} Z_m(z), \quad (16)$$

where

$$a_{nm}^{B_i} = \frac{1}{k_n^2 - k_m^2} \int_0^D Z_m(z) Z_n'(z) \left(\beta_i \frac{\beta_0'}{\beta_0} + \beta_i' \right) dz, \quad (17)$$

$$\frac{\partial k_n}{\partial B_i} = \frac{1}{2k_n} \int_0^D Z_n(z) Z_n'(z) \left(\beta_i \frac{\beta_0'}{\beta_0} + \beta_i' \right) dz, \quad (18)$$

$$\frac{\partial Z_n(z)}{\partial A_i} \approx \frac{Z_n(z) \Delta r}{k_n} \int_0^D \frac{\alpha_i(z)}{\rho(z)} k(z) Z_n^2(z) dz. \quad (19)$$

With these derivatives, we are nearly ready to carry out the inversion. Given our background mode functions, and measurements of the actual mode functions, we could move on to the actual inversion. However, in the real world, where source levels are often not known precisely, and where precise instrument calibration can be an issue, using the mode functions themselves can be problematic. Further, since the expression for the field always contains the product of the mode functions at two depths, it is usually not possible to measure the mode function by itself. Because of these issues we will actually use the ratio of each mode function to the first mode function. This eliminates source level and calibration concerns and only slightly complicates our calculation. If we define the quantity

$$m_n(z, z_s) \equiv \frac{Z_n(z) Z_n(z_s)}{Z_1(z) Z_1(z_s)}, \quad (20)$$

we can compute the derivative using the product and quotient rules, and our earlier results. Note that when using the adiabatic approximation,¹⁰ $Z_n(z)$ is computed at the receiver location, and $Z_n(z_s)$ at the source location. The derivative of the ratio with respect to some parameter γ can be written as

$$\begin{aligned} \frac{\partial m_n}{\partial \gamma} &= \frac{\partial Z_n(z_s) / \partial \gamma Z_n(z) Z_1(z_s) Z_1(z)}{Z_1^2(z_s) Z_1^2(z)} \\ &+ \frac{\partial Z_n(z) / \partial \gamma Z_n(z_s) Z_1(z_s) Z_1(z)}{Z_1^2(z_s) Z_1^2(z)} \\ &- \frac{\partial Z_1(z_s) / \partial \gamma Z_n(z_s) Z_1(z_s) Z_n(z)}{Z_1^2(z_s) Z_1^2(z)} \\ &- \frac{\partial Z_1(z) / \partial \gamma Z_n(z_s) Z_1(z_s) Z_n(z)}{Z_1^2(z_s) Z_1^2(z)}. \end{aligned} \quad (21)$$

Note that the first mode has been selected somewhat arbitrarily here. Any other mode can be used in the denominator if it is more convenient. However, the first mode is usually the best choice, since it is nonzero everywhere in the water column and is usually less affected by changes in the bottom than the other modes.

It should also be pointed out that by dealing with the mode ratios we are implicitly assuming that there is no mode coupling. This method is intended for use in cases where the adiabatic approximation¹⁰ is valid. In cases where this assumption is violated the transfer of energy between modes can give rise to apparent changes in the bottom properties which do not actually reflect the true bottom. The method can still be used as long as the true mode ratios, as defined earlier, are used as input. However, measuring those quantities when there is mode coupling is a nontrivial problem.

At this point we are ready to set up the inversion. We create a matrix, $[\partial m / \partial X]$, each row of which contains the partial derivatives of one mode function at one depth with respect to each of the parameters sought. For example, if the parameters sought are X_1, X_2, B_1, A_1 , the first row of $[\partial m / \partial X]$ is

$$\left[\frac{m_2(z_1, z_s)}{\partial X_1} \frac{m_2(z_1, z_s)}{\partial X_2} \frac{m_2(z_1, z_s)}{\partial B_1} \frac{m_2(z_1, z_s)}{\partial A_1} \right], \quad (22)$$

where z_1 is the first measurement depth. We start with m_2 because m_1 is always equal to 1. The next few rows of $[\partial m / \partial X]$ may be the same derivatives but evaluated at different measurement depths. It should be pointed out that using multiple measurement depths does not necessarily give any additional, independent information. However, by using more depths one gets more robustness to measurement noise, which is especially useful when one depth is near a null of one of the modes.

There are two other quantities that are needed before we can carry out the inversion: the column vector \mathbf{X} , which contains all the parameters we seek, and the column vector $\Delta \mathbf{m}$, which contains the difference between the mode ratios at each depth and mode number. Some might find it more intuitive to think of the \mathbf{X} vector as a $\Delta \mathbf{X}$ vector, which contains the *differences* between the desired parameters in the background model and the true case. This notation becomes a bit cumbersome, however, when partial derivatives become involved, so we have opted for calling the vector just \mathbf{X} .

The various components having been collected, we now have

$$\left[\frac{\partial m}{\partial X} \right] \mathbf{X} = \Delta \mathbf{m}. \quad (23)$$

We can solve this equation in the least-squared error sense using the pseudoinverse matrix:¹¹⁻¹³

$$\mathbf{X}_{LS} = \left[\frac{\partial \tilde{m}^\#}{\partial X} \right] \Delta \mathbf{m}, \quad (24)$$

where $[\partial \tilde{m}^\# / \partial X]$ is the pseudoinverse matrix of $[\partial m / \partial X]$.

It must be noted that we have worked so far under an approximation of a linear relationship between the variables when performing this inversion, when in reality the relation-

ship is nonlinear. Therefore we should not expect this single-step inversion to give us the correct bottom parameters. However, if our background model resembles reality, then our answer should at least give a correction to our background, which can be used to generate a new background, and the process can be repeated. After a few (usually ~ 30 or less for cases we have examined) iterations, the algorithm should converge to the values that give the best possible fit (in a least-squares sense) to the input data, given our parametrization. It must be kept in mind, however, convergence is not guaranteed. An initial background model too different from the true bottom, or a parametrization of the bottom which does not capture all important features can result in divergence.

III. SOURCES OF ERROR

To properly interpret the output of inversion algorithm, it is necessary to understand the errors that may be present in it. In this section we will discuss several sources of errors which can contribute to the overall uncertainty in the final estimate of the geoacoustic parameters. Even when the errors in the input vector are small, careless use of the pseudoinverse can lead to large errors in the estimates of the bottom parameters. Avoiding this possibility requires an understanding of how singular value decomposition is used to obtain the pseudoinverse matrix. As singular value decomposition is a well-known technique, we will not describe it here, but instead direct the reader to any standard text on matrix algebra, such as Ref. 14.

In cases where the pseudoinverse is a true inverse matrix, there are three types of errors to consider: (1) Errors in the measurement vector, (2) errors due to overly simple parametrization of the bottom, and (3) errors due to convergence to an incorrect local minimum. If the iteration process converges to the correct answer, and the parametrization of the bottom is sufficient to match reality, then the error in the estimate is linearly related to the error in the measurement. The bias of the input data should be zero. If it is not, it should be subtracted from the input data before the inversion. For completeness, however, we state that if there is bias in the input, the bias in the parameter estimation is equal to the input bias times the inversion matrix:

$$\langle \mathbf{X}_{LS} \rangle = \left[\frac{\partial \tilde{\mathbf{m}}^\#}{\partial \mathbf{X}} \right] \langle \Delta \mathbf{m} \rangle. \quad (25)$$

Similarly, the variance of the estimate is also dependent on the variance of the input:

$$\text{cov}(\mathbf{X}_{LS}) = \left[\frac{\partial \tilde{\mathbf{m}}^\#}{\partial \mathbf{X}} \right] \text{cov}(\Delta \mathbf{m}) \left[\frac{\partial \tilde{\mathbf{m}}}{\partial \mathbf{X}} \right]^T. \quad (26)$$

This expression allows us to compute the covariance of our inversion results based on the covariance of our input data, and thus quantify the uncertainty in our results.

It must be kept in mind, however, that implicit in this expression are the assumptions that the algorithm has converged to the correct minimum, that the errors in the input data are small enough that they propagate linearly through the process to the inversion results, and that the parametriza-

tion is sufficient to capture all features of the bottom. Errors in the measurement vector are beyond our control. If they are so large as to violate a linear approximation, there is little hope for the algorithm.

The other two sources of error are in competition to some degree. The more parameters used to describe the bottom, the less likely the algorithm is to miss important bottom features. However, the more parameters that are used, the more the problem becomes underdetermined. This results in a greater reliance on the smoothness and minimum norm assumptions used to compute the pseudoinverse matrix, which may not be valid.

Further understanding can be gained by examining the resolution matrix. This is especially informative when the problem is underdetermined, as is often the case in such problems. The resolution matrix, $[R]$, is defined:

$$[R] = [V][V]^T, \quad (27)$$

where $[V]$ is the matrix composed of the eigenvectors of $[\partial m / \partial X]^T [\partial m / \partial X]$. When the rank, p , is equal to the number of unknowns, N , then $[R]$ will equal the identity matrix. Otherwise, $[R]$ will be such that its elements minimize the quantity

$$\sum_i \sum_j ([R]_{ij} - \delta_{ij})^2, \quad (28)$$

where δ_{ij} is the Kronecker delta function. The degree to which $[R]$ resembles the identity matrix gives the user an idea of the resolution of the inversion result, in that off-diagonal terms will represent “smearing” of one bottom parameter into the others. If the $[R]_{ij}$ is nonzero, that indicates that a change in the i th parameter will show up also as a change in the j th parameter in the inversion result. The resolution matrix is of particular usefulness when using c -at-each- z type bottom parametrizations, because parametrizations that use a small number of parameters are less likely to be underdetermined. Backus and Gilbert¹⁵ address the concept of resolution in much greater detail.

It must also be pointed out that the pseudoinverse matrix may very well be unstable if all singular values are used to compute it. It is usually best to ignore very small singular values when computing the pseudoinverse to make the inversion more robust to error. Doing so reduces the resolution of the inversion, but will also reduce the variance in the answer. Deciding the minimum size of singular values which can be tolerated can be difficult, and some trial-and-error may be necessary to strike the right balance.

IV. EXAMPLES OF MODE AMPLITUDE INVERSION

In this section we demonstrate the feasibility of the mode amplitude perturbation algorithm by applying it to two example cases. The first is a simple synthetic data case without input errors, which serves as a basic proof of concept, while the second makes use of experimental data. For now, we keep the mode amplitude perturbation separate from the already-proven eigenvalue perturbation method in Ref. 1, so




FIG. 1. (Color online) The sound speed profile and experimental geometry for the first example.

as to make sure our new method works as expected. In the next section, we discuss how the two methods can be combined.

The first test of the method is a very simple, two-parameter, range-independent example with perfect input data. We hypothesize a “true” waveguide of 75 m of isovelocity water overlying 25 m of layered sediment. Below this, we add a half-space which continues to infinite depth. We excite a pressure field by adding an acoustic point source at 50 m depth, emitting a pure tone at 50 Hz. We assume a constant, known bottom density of 1800 kg/m^3 and zero attenuation. We use the normal mode program KRAKEN¹⁶ to generate the mode functions that will be excited in this waveguide. For this first test, we will use the mode functions directly, rather than computing the field and trying to estimate them. This separates out any potential problems in the inversion method from problems in the mode function estimation algorithm. Since we are using the mode functions directly, and the problem is range independent, we do not need to specify a range for the VLA. We do, however, need to specify the depths of the receivers, and in this case, we will use a six-element VLA with 10 m spacing. The general setup is shown in Fig. 1.

The information given so far is sufficient for solving the forward problem. However, to perform the inversion we also need to specify a background model, and parametrize the bottom. We choose a simple two-parameter bottom model which describes the bottom using only the sound speed at the water-bottom interface, and the sound speed gradient. This model is appealing because the low number of parameters allows us to solve a fully determined problem as long as we have at least three propagating modes, and because it is often capable of describing the true bottom as a reasonable first approximation. While there are layers in the true bottom, the rate of increase of the sound speed of the layers is nearly linear with depth, so a gradient can do well to approximate the sound speed profile.

For the background model, we use a Pekeris-like waveguide with a sound speed of 1650 m/s throughout the bot-

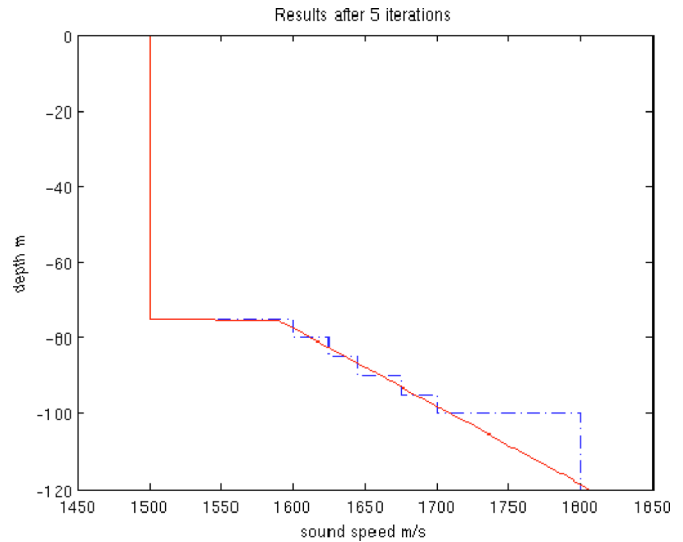


FIG. 2. (Color online) The results of the geoacoustic inversion algorithm after five iterations. The dotted curve shows the “true” sound speed profile, while the solid curve indicates the inversion result.

tom. However, unlike the standard Pekeris model, we include a pressure-release surface at 125 m depth. The reason for this is that the expression for the derivative of the mode function includes a sum over all modes. In order to have a complete set, we must have an infinite number of modes, which is only possible if we include the nonpropagating modes that have complex eigenvalues. Solving for complex eigenvalues is a difficult task, and greatly increases the computation time needed for solving the problem. By adding a pressure-release false bottom, we can maintain a proper Sturm–Liouville problem, and the eigenvalues of the nonpropagating modes become purely imaginary, and thus much easier to compute. The error introduced by our use of this false bottom should be small if it is placed at a depth below the turning point of the highest propagating mode.¹⁷

With all of the necessary components in place, we can now execute the inversion process. The mode functions are computed for the background model, and the mode ratios computed for each mode, and each receiver depth. These are compared to the mode ratios formed using the output from KRAKEN. The derivatives of the ratios are computed following the results of Sec. II. The linear set of equations is solved using the pseudoinverse, giving the necessary changes in the two parameters. These changes in the parameters are incorporated into the background model, and the process repeated again. After five iterations of this process, the algorithm converges to the result shown in Fig. 2, which agrees well with the true profile.

There is nontrivial disagreement below 100 m depth, but using our simple two-parameter model, this is unavoidable. For the first 25 m of the bottom, the fit is as good as our two-parameter model allows, showing that given perfect input data, the method will provide a good estimate of the bottom sound speed profile (SSP). Also, since this result was obtained using a background model containing a pressure-release false bottom, we gain confidence that using such false bottoms in a proper fashion in our background model will not corrupt the inversion results.




FIG. 3. (Color online) The pressure magnitude vs range measured during MOMAX 2, at 175 Hz.

For our second example we use data from the LWAD 99-1/MOMAX 2 Experiment, which was carried out in the Gulf of Mexico in February 1999.⁴ During this experiment the mode amplitudes were estimated by taking a Hankel transform of pressure versus range data, a procedure described in Ref. 9 and used in Ref. 1. A single hydrophone suspended from a drifting buoy moved through the sound pressure field created by a single 175 Hz pure tone source, suspended from the moored R/V Gyre. A synthetic aperture HLA was formed as the receiver drifted, providing the input data for the Hankel transform. This was possible because the waveguide was very nearly range independent over the section of data used, and the 70-m-deep water column was stable during the time period over which the synthetic aperture was formed.

Figure 3 shows the pressure magnitude as a function of range measured by the drifting receiver, after the raw pressure measurement was demodulated for the source frequency of 175 Hz. The magnitude and phase data were then Hankel transformed to obtain an estimate of the depth-dependent Green's function versus horizontal wave number, which is shown in Fig. 4.

The 19 modal peaks used as input data to the inversion algorithm are indicated with asterisks. It may seem unreasonable to assume that all of these peaks are actual modal peaks rather than sidelobes of the finite-aperture transform, but each peak lined up surprisingly well with the expected eigenvalue locations based on prior estimates¹⁸ of the geoacoustic parameters for this location. Further, the widths of the peaks also seem to indicate main lobes rather than sidelobes, as can be seen by observing the narrow sidelobes to the right of the first (largest eigenvalue) mode. It is possible that some of the modes have been misidentified, but we assume the large majority of the modal peaks are correctly interpreted. The possibility of misidentifying modal peaks is a common problem when dealing with experimental data, and is a problem any modal inversion method must be able to handle.

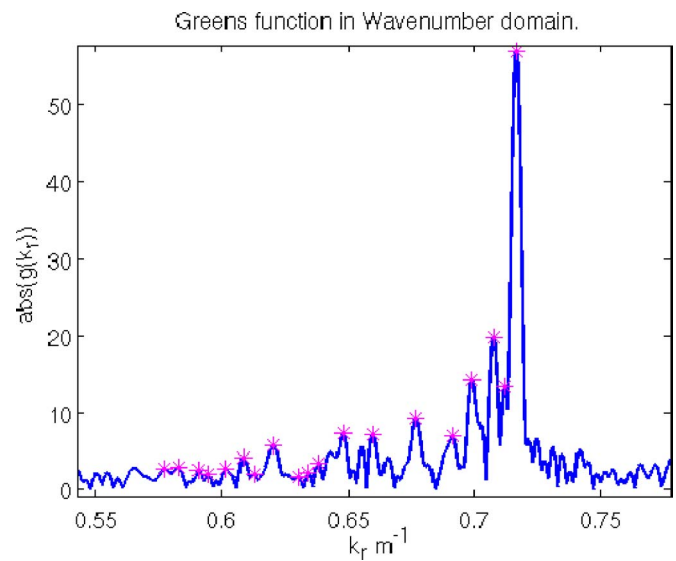


FIG. 4. (Color online) The Hankel transform of the pressure vs range data shown in Fig. 3. Mode peaks used in the inversion are indicated with asterisks.

In order to reduce the errors introduced by a few misinterpreted peaks, we use a simple two-parameter model for the sea floor. By solving an overconstrained problem we tend to reject unbiased errors in our input, whereas if we solve an underdetermined problem the algorithm will introduce erroneous bottom properties in order to fit the errors introduced by the measurement. The cost of this, however, is that we are limited to inverting very simple sound speed profiles in the bottom. Fortunately, the pre-cruise estimate in Ref. 18 predicts a bottom that could be described reasonably well with a two-parameter model. If this were not the case, we would have to use more parameters to describe the bottom.

For a background model, we again choose a simple Pekeris-like model. While we have a better *a priori* estimate of the bottom SSP in this case, one usually will not have such a good starting point, so it is a better test of the method to start with a less accurate background model. We use our background model to compute the mode ratios at the receiver depth, and compare those to our measurements. We compute the derivatives of the ratios with respect to our two parameters, and solve the linear system for the necessary changes in the parameters, as described in Sec. III. We repeat this process until we have convergence, and the results are shown in Fig. 5.

It is clear that the final inversion result matches the pre-cruise estimate well. That 25 iterations were required before convergence may be a sign that some modes were misidentified, or that our bottom parametrization failed to capture some of the features of the true sound speed profile. It is also possible that some of the parameters which were treated as known, such as the density profile, the attenuation, and the water column SSP, were not accurate. Despite all these possibilities, however, we achieve a result which compares well with our expectation, so we can be confident the mode amplitude perturbation algorithm works.




FIG. 5. (Color online) The results of the geoacoustic inversion of the 175 Hz data from MOMAX 2. The vertical dashed line shows the starting model of the sound speed profile, and the solid curves show iterations 1, 5, 10, and 15. The bold solid curve is the final result after convergence at 25 iterations. The step-like dashed curve is the best estimate of the true bottom profile available before the experiment.

V. COMBINING MODE AMPLITUDE AND EIGENVALUE PERTURBATION METHODS

Having seen that mode amplitude perturbation alone can produce a suitable inversion, we now compare the quality of the inversion it produces to that of eigenvalue perturbation method, and then consider the possibility of combining the two methods to provide a better result than either produces on their own. To compare the two methods, we require some metric of performance. The two most common metrics of performance of an inversion algorithm are the bias and variance of the error in the inverted parameters.

As stated earlier, assuming the algorithm converges to the correct minimum, and that the measurement errors are small enough that they affect the estimate linearly, we can compute the bias and variance of our inversion from the bias and variance of the input data using Eqs. (25) and (26). With these expressions for error covariance in hand, it is possible to compare the performance of the two inversion methods. The “better” algorithm is the one with the lower error covariance. However, it is often the case that one method will perform better at estimating one parameter, while the other performs better at estimating another.

For example, eigenvalue perturbation tends to provide a lower-variance estimate of subbottom features, whereas mode amplitude perturbation methods do better at estimating shallower parameters such as interface speeds.¹⁹ This is not unexpected, since the mode functions contain little energy at large depths, and thus are not much affected by the subbottom, whereas the eigenvalues are determined by the boundary conditions, and thus are highly affected by subbottom parameters. Thus, which method is preferred will depend, in part, on which parameters the user is most concerned with estimating accurately.

Further, which method is better will depend on the quality of the input data. If the eigenvalues are estimated very

precisely, then eigenvalue perturbation will be superior. On the other hand, if the mode amplitudes are estimated well, but the eigenvalues have significant error, mode amplitude perturbation should be used. Which inversion method is more effective will also depend somewhat on the environment, as this will change the inversion matrix. In short, in order to determine the preferred method for a given scenario, one must perform the inversions and compare the computed error variances.

Since both methods can provide lower-variance estimates for different parameters, it is logical to ask if we can combine the two methods for even better estimates. At first it might seem best to simply take the values of each parameter from the inversion with the lower error variance. However, the estimates of the various parameters are not independent, and mixing and matching them can lead to problems. A better solution is to use all the information available to both methods at once. To do this we can combine the two derivative matrices into one, and the two data vectors into one, forming a new equation:

$$\begin{bmatrix} \frac{\partial Q_{MA}}{\partial X} \\ - \\ \frac{\partial Q_E}{\partial X} \end{bmatrix} \mathbf{X} = \begin{bmatrix} \Delta Q_{MA} \\ - \\ \Delta Q_E \end{bmatrix}. \quad (29)$$

This equation is essentially the same as the one used by either of the methods, but we cannot solve it in quite the same way (i.e., using the Morse–Penrose pseudoinverse), because the two types of input data differ significantly in size. Eigenvalue perturbations are typically on the order of 0.001, whereas mode ratio amplitude perturbations are closer to order 0.1 or even 1. Using the pseudoinverse to find a least-squares solution to Eq. (29) would essentially result in the same answer as if we just used the mode amplitude equation. The eigenvalue perturbations are just too small to affect the answer significantly. Thus, we must take into account the relative sizes of the perturbations rather than the absolute sizes. One way of doing this is to weight each equation (i.e., multiply each row of the above-presented matrix and the corresponding entry in the data vector) by the standard deviation of the perturbation, and then compute the pseudoinverse.

An even better method is to make use of the stochastic inverse (SI).^{20,21} This method for creating the inverse matrix not only accounts for the relative sizes of the different perturbations but also for the relative sizes of the changes in the parameters. For example, an interface speed can easily change by 10 m/s over an experiment, whereas changes in a gradient are likely to be much smaller. We will leave the derivation of the stochastic inverse to the references, and state here only the expression for the inverse. For

$$\mathbf{d} = [G]\mathbf{q} + \mathbf{e}, \quad (30)$$

where \mathbf{d} is the measurement vector, \mathbf{e} a zero-mean noise vector with covariance R_e , and \mathbf{q} the unknown, zero-mean parameter vector with covariance R_q , the stochastic inverse solution is

$$\hat{q}_{SI} = [\tilde{G}^{SI}]d = [R_q][G]^T[[G][R_q][G]^T + [R_e]]^{-1}d. \quad (31)$$

This inversion minimizes the errors squared, but normalizes by the uncertainty in the quantity with the error. For example, an error in a quantity with variance 0.0001 should be much less than an error in a quantity with a variance of 0.1. The inversion also tries to minimize the q vector, again weighted by the variance of the elements in the vector.

Though the stochastic inverse requires that the user supply considerably more information, which may have to be estimated, it allows one to combine the eigenvalue and mode amplitude perturbation results and find a solution that takes into account all the data available in the pressure field. By looking at the covariance of our estimate, we can determine which of the three methods (mode amplitude, eigenvalue, or the combination) gives the best estimate of the bottom parameters.

Another advantage of combining the two methods is a decreased chance of converging to an incorrect local minimum. While there is little danger of this happening when only a small set of parameters is used to describe the bottom, it is quite possible when a large parameter set is used. The more input data there are, the more “ways out” of the local minima there are, and the less likely the algorithm is to get trapped. The possibility of convergence to an incorrect local minimum is a problem that the user must keep in mind. If the algorithm converges to a profile that does not seem realistic, it is likely that it has reached a local minimum, and the user should try starting with a different background model, to see if the same result is obtained.

Even if the profile produced does look reasonable and reproduces the measurements well, it is still possible (though unlikely) that the profile does not match the true bottom! Unfortunately, this is inescapable due to the underdetermined nature of the problem. That said, using more input data does tend to reduce the risk of local minima. This is only true, of course, if the additional input data contain at least some independent information. While much of the information contained in the mode amplitudes is redundant to that in the eigenvalues, some of the information is independent. Using both sources of information should reduce the chances of there being reasonable-looking profiles at incorrect local minima.

In order to illustrate the combined eigenvalue and mode amplitude perturbation method we examine a synthetic example. For this case we will posit a waveguide 200 m deep and 4000 m in length. The water column is 50 m deep and is treated as isovelocity (1500 m/s, 1000 kg/m³). Below this is a 70 m deep, range-dependent sediment layer. Two parameters (interface speed and gradient) are sufficient to describe this layer at any given range. Density is treated as constant in the layer (1600 kg/m³). Below the sediment layer is a range-independent subbottom (1800 m/s, 1800 kg/m³). This waveguide is shown in Fig. 6.

A synthetic field was generated for the waveguide, simulating a VLA at range 0, and a 100 Hz source moving from 200 to 4000 m from the VLA, at a constant depth of 30 m. Seven propagating modes were used to generate the synthetic field. The VLA had nine elements, with 5 m spacing,




FIG. 6. (Color online) Sound speed contours of the synthetic waveguide. The contours indicate lines of constant sound speed. The water column and subbottom are treated as isovelocity and range independent, whereas the sound speed in the sediment layer varies both in depth and range.

going from 5 m depth to 45 m. Zero mean, Gaussian random noise was added to the field. The modal parameters were then estimated from the VLA measurements. The method by which this was done will be described in a future paper, and it is outside the scope of this paper. All that is necessary for the inversion is that we have an estimate of the parameters, and an estimate of the covariance of our parameter estimates. Figures 7 and 8 show the estimated and true values of the modal parameters as a function of range.

With the estimates of the mode amplitudes and eigenvalues and their variances in hand, we can estimate the bottom parameters. For the inversions we treat the water and subbottom as known, and invert for the interface speed and gradient of the sediment layer. We also treat density within the sediment layer as known and ignore attenuation. As with the first




FIG. 7. (Color online) Eigenvalues estimates used to test the combination of eigenvalue and amplitude perturbation methods. The values of the seven eigenvalues at the source location as a function of range are shown. The solid lines indicate the true values, while the estimates are shown with dotted lines.




FIG. 8. (Color online) Mode amplitude estimates used to test the combination of eigenvalue and amplitude perturbation methods. The mode amplitudes at 25 m depth on the VLA as a function of source-receiver range are shown. The true values are depicted with solid lines, and estimated values are shown as dotted lines.

test case, this is done to keep our example simple, and is not a requirement of the method. In order to compare the mode amplitude perturbation method with the eigenvalue perturbation method we do separate inversions using each, and compute the error variances using Eq. (26).

The error variance of the four estimated parameters is shown in Fig. 9. In this case, the error variances favor the mode amplitude perturbation for some of the parameters and eigenvalue perturbation for others. For example, the interface speed at the source location as estimated by mode amplitude perturbation has a lower error variance than it does when estimated by eigenvalue perturbation. But the opposite is true

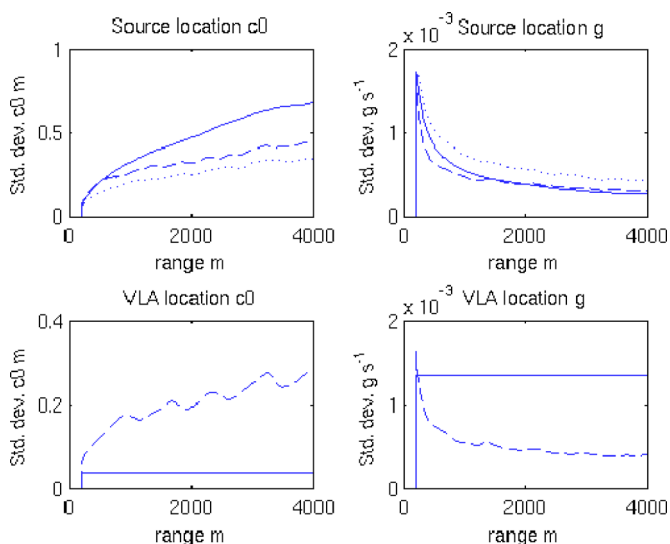


FIG. 9. (Color online) Error variance of the four inverted parameters. The eigenvalue perturbation result is shown as solid lines, the mode amplitude result in dashed lines, and the combined method shown in dotted lines. For the bottom two plots, the curve for the combined method overlays that of the eigenvalue method.




FIG. 10. Interface sound speed at the source, real and inverted. The true interface speed (solid) as a function of range for the synthetic experiment is shown. The three inverted values are also shown. The eigenvalue perturbation result is shown as dashed lines, the mode amplitude inversion as dotted lines, and the combination method result in dash-dotted lines. For most of the experiment, the combination method tends to track whichever of the two inversion results is closer to the true value.

of the receiver-position estimate of the interface speed. Estimated error variance of the source-position gradient is similar with both methods, but the VLA-position estimate of the gradient appears to favor mode amplitude perturbation. This, however, is an error on the part of the parameter estimation algorithm, which has supplied an overconfident estimate of the variance. It is informative to track how such an underestimate of the variance in our estimate of a parameter propagates through our algorithm and affects the inversion result.

Using the stochastic inverse, the estimated covariance of the errors in our measured acoustic parameters, and an estimated matrix of the unknown parameter covariance (the parameters were considered uncorrelated, but the true variances of each parameter were used), we can solve the combined mode amplitude/eigenvalue perturbation equation. And just as we did with the pseudoinverse, we can use the stochastic inverse to compute the error variance of our inversion. When we do so for our synthetic example, we get the best of both methods, which is shown in Fig. 9. As can be seen, the estimate of the source location interface speed is improved over either method, and the variance of the estimate of the VLA position parameters matches that of the eigenvalue method. On the other hand, there is a slight decrease in the quality of the estimate of the gradient at the source position. This is due to our overconfidence in the mode-amplitude method's estimate of this parameter.

Now that we have looked at the predicted accuracy of our inversion, let us examine the actual accuracy. Figures 10–13 show the true values of the parameters along with the inverted values from the three methods. Figure 10 shows the inversion results for the interface speed at the source position. While all three methods are fairly comparable in quality, for most of the experiment the combination result tracks the better of the two other inversions. The differences are not




FIG. 11. Interface speed at the VLA, real and inverted. The eigenvalue perturbation result overlays the constant true value of the VLA-position interface speed. As the estimate of the VLA position eigenvalue changes very little, the inverted value of the parameter changes very little as well. The mode amplitudes, however, can change due to changes at the source, or at the VLA, and thus the mode amplitude perturbation result varies with range, and has much larger error. The combination inversion method tracks the result of the eigenvalue perturbation well, and thus remains at the true value.

particularly dramatic in this case, but over the course of the experiment, the combination result does appear to be the best of the three.

Figure 11 shows the inverted values of the interface speed at the VLA position. Here the improvements of the combination method are clear. The eigenvalue perturbation

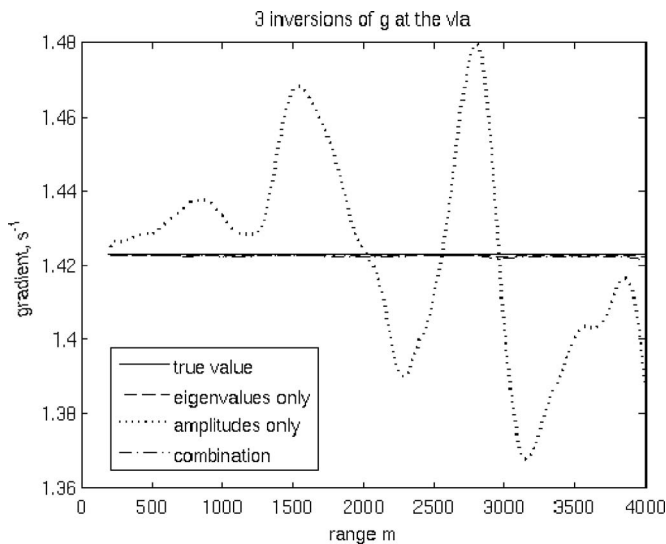


FIG. 12. Gradient at the VLA, real and inverted. As in Fig. 11, the true value is constant, and the unchanging local eigenvalue estimate gives a good result for the eigenvalue perturbation method. The mode amplitude perturbation method, however, has trouble distinguishing between local and source-position changes, and thus provides a poorer estimate of the parameter. The combination result tracks the eigenvalue perturbation result, and the true value, well. Note that the level of error in the mode amplitude perturbation result is at odds with the predicted error variance for this parameter shown in Fig. 9. This is due to an overconfidence (i.e., too low a value for the estimated variance) in some of the estimates used as input data.




FIG. 13. Gradient at the source, real and inverted. The source-position gradient as a function of source-receiver range is shown. The true value is shown with a solid line. All three estimates are comparable, and show similar amounts of error. However, the combination method does the best at tracking the features of the parameter with range.

method, using an estimate of the unchanging local eigenvalues, does a good job of tracking the true value of the parameter. The mode amplitude method, however, uses mode ratios which are affected by changes at both the source and receiver position, and because of this struggles to distinguish between changes at the source, and changes at the VLA. Thus, its estimate of the parameter is nonconstant, and has significantly larger errors than the eigenvalue perturbation method. The combination method is able to use the local eigenvalue estimate, and thus tracks the true value well. Figure 12 is similar. Again the mode amplitude perturbation method struggles to distinguish between local and source-position changes, and thus gives a nonconstant estimate with larger errors. The methods that make use of the estimate of the local eigenvalue, however, do a good job at tracking the parameter. Note that Fig. 9 would lead us to believe that the mode amplitude estimate would be more precise for this parameter, and yet clearly the eigenvalue estimate is superior. The combined method tracks the eigenvalue result, as we would hope, despite the overconfidence of our variance estimate.

Figure 13 shows the estimates of the gradient at the source position. Similar to the source-position interface speed estimates, all three methods have similar levels of bias, though the mode amplitude and combination methods provide better estimates of this parameter over the last kilometer. Additionally, the combination method does significantly better than the other two methods at tracking the features of the parameter in range. The maxima and minima of the parameter are clearly visible in the combination method estimate, though the values of the estimate are slightly off.

Overall, while the differences between the various methods were not drastic, the combination method tended to do slightly better than the other two methods. Its results tended to resemble those of whichever method produced the better estimate of each parameter. As stated earlier, which method

is best will depend on the quality of the estimates of the input data and the parametrization of the bottom selected. Since the combination method tends to track the better of the two estimates in each case, it can be expected to provide results similar to the better method, whichever method that happens to be in a give case.

VI. CONCLUSIONS

In this paper we have derived the results necessary to carry out a mode amplitude perturbative inversion method for determining the geoacoustic properties of the seabed. We have applied this algorithm to two data sets: One a simple synthetic data case with a known solution, and the other a LWAD 99-1 experimental data set, where a previous estimate of the bottom sound speed profile was available. In both cases the algorithm was able to successfully estimate the bottom geoacoustic parameters. The modal amplitude inversion algorithm is analogous to the modal eigenvalue inversion algorithm described in Ref. 1, the major difference being the type of input data that is used. As a consequence, it is natural to compare the two methods and to combine them so as to gain the benefits of each. Using the stochastic inverse matrix approach, we demonstrated the effectiveness of this hybrid inversion algorithm.

ACKNOWLEDGMENTS

Funding for the research presented here was provided by the Office of Naval Research, and the WHOI Academic Programs Office.

¹S. D. Rajan, J. F. Lynch, and G. V. Frisk, "Perturbative inversion methods for obtaining bottom geoacoustic parameters in shallow water," *J. Acoust. Soc. Am.* **82**, 998–1017 (1987).

²T. B. Neilsen and E. K. Westwood, "Extraction of acoustic normal mode depth functions using vertical line array data," *J. Acoust. Soc. Am.* **111**, 748–756 (2002).

³P. Hursky, W. S. Hodgkiss, and W. A. Kuperman, "Matched field processing with data-derived modes," *J. Acoust. Soc. Am.* **109**, 1355–1366 (2001).

⁴G. V. Frisk, K. M. Becker, and J. A. Doust, "Modal mapping in shallow water using synthetic aperture horizontal arrays," *OCEANS 2000 MTS/IEEE Conference and Exhibition*, 11–14 September 2000, Providence, RI, Vol. 1, pp. 185–188.

⁵C. T. Tindle, L. M. O'Driscoll, and C. J. Higham, "Coupled mode perturbation theory of range dependence," *J. Acoust. Soc. Am.* **108**, 76–83 (2000).

⁶G. V. Frisk, *Ocean and Seabed Acoustics: A Theory of Wave Propagation* (Prentice Hall, Englewood Cliffs, NJ, 1994).

⁷F. B. Jensen, W. A. Kuperman, M. B. Porter, and H. Schmidt, *Computational Ocean Acoustics* (AIP Press, New York, 1994).

⁸C. L. Pekeris, "Theory of propagation of explosive sound in shallow water," *Mem.-Geol. Soc. Am.* **27**, 1–117 (1948).

⁹S. D. Rajan, G. V. Frisk, and J. F. Lynch, "On the determination of modal attenuation coefficients and compressional wave attenuation profiles in a range-dependent environment in Nantucket sound," *IEEE J. Ocean. Eng.* **17**, 118–128 (1992).

¹⁰A. D. Pierce, "Extension of the method of normal modes to sound propagation in an almost-stratified medium," *J. Acoust. Soc. Am.* **37**, 19–27 (1965).

¹¹E. H. Moore, "On the reciprocal of the general algebraic matrix," *Bull. Am. Math. Soc.* **26**, 394–395 (1920).

¹²R. Penrose, "A generalized inverse for matrices," *Proc. Cambridge Philos. Soc.* **51**, 406–413 (1955).

¹³R. Penrose, "On best approximate solution of linear matrix equations," *Proc. Cambridge Philos. Soc.* **52**, 17–19 (1956).

¹⁴D. C. Clay, *Linear Algebra and its Applications* (Addison-Wesley, New York, 1994).

¹⁵G. Backus and F. Gilbert, "Uniqueness in the inversion of inaccurate gross earth data," *Proc. R. Soc. London, Ser. A* **266**, 187–269 (1970).

¹⁶M. B. Porter, "The KRAKEN normal mode program," Technical Report, SACLANT Undersea Research Center, La Spezia, Italy, 1991.

¹⁷A. O. Williams, Jr., "Hidden depths: Acceptable ignorance about ocean bottoms," *J. Acoust. Soc. Am.* **59**, 1175–1179 (1976).

¹⁸B. R. Gomes, J. K. Fulford, and R. Nero, "Preliminary pre-experiment environmental characterization for the Littoral Warfare Advanced Development (LWAD) 99-1 SE Gulf of Mexico," Naval Research Laboratory, Stennis Space Center, MS, 1998.

¹⁹T. L. Poole, "Geoacoustic inversion by mode amplitude perturbation," Ph.D. thesis, MIT/WHOI Joint Graduate Program in Oceanography/Oceanographic Engineering, Cambridge and Woods Hole, MA, 2007.

²⁰J. N. Franklin, "Well-posed stochastic extensions of ill-posed linear problems," *J. Math. Anal. Appl.* **31**, 682–716 (1970).

²¹L. L. Souza, "Inversion for subbottom sound velocity profiles in the deep and shallow ocean," Ph.D. thesis, MIT/WHOI Joint Graduate Program in Oceanography/Oceanographic Engineering, Cambridge and Woods Hole, MA, 2005.