Geophysical Prospecting, 2002, **50**, 577–588

Iterative resolution estimation in least-squares Kirchhoff migration

Sergey Fomel,[1]* James G. Berryman,[2] Robert G. Clapp[3] and Marie Prucha[3]

[1]*Lawrence Berkeley National Laboratory, 1 Cyclotron Road, Mail Stop 50A-1148, Berkeley, CA 94720,* [2]*University of California, Lawrence Livermore National Laboratory, PO Box 808 L-200, Livermore, CA 94551-9900, and* [3]*Stanford Exploration Project, Mitchell Bldg., Department of Geophysics, Stanford University, Stanford, CA 94305-2215, USA*

Received February 2001, revision accepted June 2002

ABSTRACT

We apply iterative resolution estimation to least-squares Kirchhoff migration. Reviewing the theory of iterative optimization uncovers the common origin of different optimization methods. This allows us to reformulate the pseudo-inverse, model resolution and data resolution operators in terms of effective iterative estimates. When applied to Kirchhoff migration, plots of the diagonal of the model resolution matrix reveal low illumination areas on seismic images and provide information about image uncertainties. Synthetic and real data examples illustrate the proposed technique and confirm the theoretical expectations.

INTRODUCTION

Kirchhoff prestack depth migration remains the most widely used method for seismic imaging in complex areas. The method is especially attractive for 3D imaging because of its ability to handle naturally irregular acquisition geometries. The negative effect of irregular sampling on seismic images can be additionally balanced by applying the least-squares migration approach (Cole and Karrenbach 1992), which has recently attracted a lot of attention in the geophysical literature (Chavent and Plessix 1999; Duquet and Marfurt 1999; Nemeth, Wu and Schuster 1999).

According to the least-squares approach, the migration operator is constructed as a least-squares inverse of forward Kirchhoff modelling (Tarantola 1987). The inverse operator can be effectively approximated through an application of an iterative optimization scheme. The conventional migration is then considered as the adjoint of the modelling operator or, in other words, the first step of an iterative optimization process (Claerbout 1992). A more accurate representation (i.e. additional steps) can compensate for irregularities and artefacts of irregular acquisition (Nemeth 1996; Nemeth *et al.* 1999).

A blind least-squares approach cannot, however, compensate for lack of information in the input data. For example, if a particular area in the subsurface is not illuminated by reflected waves, a proper image of that area cannot be resolved by least-squares migration alone. In this case, part of the image will belong to the null space of the least-squares inverse problem. Spotting low-illumination areas is important both for making acquisition decisions and for evaluating the uncertainty of the existing images. Duquet, Marfurt and Dellinger (1998) proposed using the inverse diagonal of the Hessian matrix as a measure of illumination in Kirchhoff imaging. Although this measure does provide useful information about the problem's well-posedness, a more rigorous approach to the solution uncertainty would be to estimate the corresponding model resolution operator (Jackson 1972). The least-squares approach to resolution estimation is an alternative to the commonly used ray-tracing method for illumination analysis (Bear *et al.* 1999; Muerdter and Ratcliff 2000; Laurain and Vinje 2001).

In this paper, we show how to obtain a resolution estimate using iterative methods of linear unconstrained optimization, such as the method of conjugate gradients (Lanczos 1950; Hestenes and Stiefel 1952) and LSQR (Paige and Saunders 1982). For a given forward modelling operator, predicting the existing data from an unknown model, iterations approach the model, which minimizes the squared residual error of prediction. In linear problems, the global minimum does exist. However, finding it requires, in general, the number of iterative steps to be equal to the number of unknown model parameters. In large-scale problems, typical

*Currently at: Bureau of Economic Geology, University of Texas at Austin, University Station, Box X, Austin, TX 78713, USA.
E-mail: sergey.fomel@beg.utexas.edu

in geophysical applications, the computational cost makes complete solution practically infeasible. Nevertheless, iterative methods allow us to get a reasonable *estimate* of the solution in a small number of iterations.

When the exact *solution* of an inverse problem is replaced by an *estimate* of the solution, the inversion theory needs to be reformulated. Methods and formulations, designed for the complete solutions, are no longer applicable in the case of iterative estimates. This conclusion applies to such objects as a pseudo-inverse operator, model resolution and data resolution. These objects are conventionally associated with singular value decomposition (SVD), which becomes infeasible in large-scale problems. Berryman (2000a) reviewed different methods of iterative optimization, primarily the method of conjugate directions and the LSQR method. He proved that these methods have a common origin in the general principle of the iterative residual minimization. The general principle leads to remarkable orthogonalization properties for particular sets of vectors in the model and data subspaces.

Following the results of Berryman (2000b), we show how to define the effective pseudo-inverse operator, model and data resolution for iterative methods. Since the exact solution is not available, these definitions apply to effective iterative estimates of the corresponding operators, which were strictly defined in the inversion theory. The iterative estimates are constructed from the vectors, already appearing in the conjugate-gradient iteration. Therefore, they require minimal additional computation with respect to an iterative least-squares inversion. The diagonal of the resolution matrix can serve as a rough direct estimate of the model uncertainty (Wang and Pratt 1997). A similar, although less efficient, approach was proposed by Minkoff (1996) and Yao, Roberts and Tryggvason (1999), who applied it in conjunction with the LSQR method (Paige and Saunders 1982).

Finally, we apply the iterative technique for resolution estimation in Kirchhoff imaging. Synthetic and real data tests show that a resolution estimate can indeed provide valuable information about the uncertainty of Kirchhoff images and reveal image areas with illumination problems.

Here, we have not found any use for the data resolution estimates. In other applications, the data resolution matrix can serve for assessing inconsistencies in the data.

RESOLUTION OPERATORS FOR BOTH MODEL AND DATA

Following the theoretical discussion of Berryman (2000b) and the known properties of the iterative optimization method, summarized in the Appendix, we formulate the iterative estimates for the pseudo-inverse model resolution and data resolution operators. We use a notation in which the linear inversion problem to be solved takes the form

$$\mathbf{M}\mathbf{x} = \mathbf{d}, \tag{1}$$

where we assume that the data vector $\mathbf{d}$ and the linear forward modelling operator $\mathbf{M}$ are given and that the model vector $\mathbf{x}$ is being sought.

In the case of least-squares migration, $\mathbf{x}$ corresponds to the reflectivity model, $\mathbf{d}$ is the observed data and $\mathbf{M}$ is the Kirchhoff modelling operator (the adjoint of migration). We generally assume that the problem is overdetermined so that the number of data values exceeds the dimension of the model space.

Linear iteration

We want to solve the problem (1) in an iterative fashion, so we assume that the updates to the solution take the general form,

$$\mathbf{x}_n = \mathbf{x}_{n-1} + \alpha_n \mathbf{p}_{n-1}, \tag{2}$$

where $\mathbf{x}_{n-1}$ is the preceding estimate of $\mathbf{x}$, $\mathbf{x}_n$ is the new estimate of $\mathbf{x}$, $\mathbf{p}_{n-1}$ is some direction to be specified in the model space and α_n is an optimization parameter (or direction weight factor). Defining the residual data error as $\mathbf{r}_n \equiv \mathbf{d} - \mathbf{M}\mathbf{x}_n$, we find the general relation

$$\mathbf{r}_n = \mathbf{r}_{n-1} - \alpha_n \mathbf{M}\mathbf{p}_{n-1}. \tag{3}$$

Pseudo-inverse estimate

From (2), it follows easily that the model estimate at the nth iteration must be of the form,

$$\mathbf{x}_n = \sum_{i=1}^{n-1} \alpha_{i+1} \mathbf{p}_i, \tag{4}$$

where we assume for simplicity that $\mathbf{x}_1 = 0$. Then substituting (A2) from the Appendix – or more directly the first ratio in (A12) – for the α_i's shows that the nth iterate is given explicitly by

$$\mathbf{x}_n = \sum_{j=1}^{n-1} \frac{\mathbf{p}_j \mathbf{p}_j^{\mathrm{T}}}{(\mathbf{p}_j, \mathbf{M}^{\mathrm{T}}\mathbf{M}\mathbf{p}_j)} \mathbf{g}_j = \sum_{j=1}^{n-1} \frac{\mathbf{p}_j \mathbf{p}_j^{\mathrm{T}}}{(\mathbf{p}_j, \mathbf{M}^{\mathrm{T}}\mathbf{M}\mathbf{p}_j)} \mathbf{M}^{\mathrm{T}}\mathbf{d} \tag{5}$$

for this scheme. The resulting approximate inverse operator is therefore

$$\left(\mathbf{M}^{\mathrm{T}}\mathbf{M}\right)^{\dagger} \simeq \sum_{j=1}^{n-1} \frac{\mathbf{p}_j \mathbf{p}_j^{\mathrm{T}}}{(\mathbf{p}_j, \mathbf{M}^{\mathrm{T}}\mathbf{M}\mathbf{p}_j)}, \tag{6}$$

which form we now want to study. We use the dagger notation to indicate that the expression in (6) approximates a pseudo-inverse, because it may happen that the normal matrix is singular, in which case the standard inverse does not exist.

Although (6) might appear to be in the form of a singular value decomposition, it definitely is not. The $\mathbf{p}_n$'s are not orthogonal and the denominators of these terms are not eigenvalues. If we define the matrix composed of direction vectors at the nth iteration as

$$\mathbf{P}_n = (\mathbf{p}_1 \quad \mathbf{p}_2 \quad \cdots \quad \mathbf{p}_n), \tag{7}$$

then the approximate inverse operator can be rewritten as

$$\left(\mathbf{M}^{\mathrm{T}}\mathbf{M}\right)^{\dagger} \simeq \mathbf{P}_n \mathbf{D}_{\mathbf{P}}{}^{-1} \mathbf{P}_n^{\mathrm{T}}, \tag{8}$$

where the matrix $\mathbf{D}_{\mathbf{P}}$ is a diagonal matrix whose diagonal elements are given by $D_{jj} = (\mathbf{p}_j, \mathbf{M}^{\mathrm{T}}\mathbf{M}\mathbf{p}_j)$. In fact the entire matrix is given directly by

$$\mathbf{D}_{\mathbf{P}} \equiv \mathbf{P}_n^{\mathrm{T}} \mathbf{M}^{\mathrm{T}} \mathbf{M} \mathbf{P}_n, \tag{9}$$

because of the conjugacy of the $\mathbf{p}$'s composing $\mathbf{P}_n$. Now (A19) shows that

$$\mathbf{P}_n \mathbf{B}_n = \mathbf{G}_n, \tag{10}$$

where

$$\mathbf{G}_n = (\mathbf{g}_1 \quad \mathbf{g}_2 \quad \cdots \quad \mathbf{g}_n), \tag{11}$$

and the matrix $\mathbf{B}_n$ is bidiagonal with units along the main diagonal and β's along the upper diagonal:

$$\mathbf{B}_n = \begin{bmatrix} 1 & \beta_2^{(1)} & 0 & \cdots & 0 \\ 0 & 1 & \beta_3^{(2)} & \cdots & \cdots \\ 0 & 0 & \cdots & \cdots & 0 \\ \cdots & \cdots & \cdots & 1 & \beta_n^{(n-1)} \\ 0 & \cdots & 0 & 0 & 1 \end{bmatrix}. \tag{12}$$

Multiplying (10) on the right by the inverse of $\mathbf{B}_n$ and then substituting in (8), we find that

$$\left(\mathbf{M}^{\mathrm{T}}\mathbf{M}\right)^{\dagger} \simeq \mathbf{G}_n \mathbf{B}_n^{-1} \mathbf{D}_{\mathbf{P}}{}^{-1} \left(\mathbf{B}_n^{\mathrm{T}}\right)^{-1} \mathbf{G}_n^{\mathrm{T}}. \tag{13}$$

Thus, the approximate inverse is seen to have the general form,

$$\left(\mathbf{M}^{\mathrm{T}}\mathbf{M}\right)^{\dagger} \simeq \mathbf{G}_n \mathbf{T}_n^{-1} \mathbf{G}_n^{\mathrm{T}}, \tag{14}$$

where $\mathbf{T}_n$ is the tridiagonal matrix,

$$\mathbf{T}_n = \mathbf{B}_n^{\mathrm{T}} \mathbf{D}_{\mathbf{P}} \mathbf{B}_n. \tag{15}$$

This result highlights the similarities between the conjugate-gradient method and other iterative methods such as that of Lanczos (1950) and LSQR (Paige and Saunders 1982), which also produce tridiagonal representations of the matrix to be inverted.

Model resolution estimate

Although the tridiagonal form found in (15) is interesting in its own right, the more important result contained in (14) is the fact that this analysis has resulted in a decomposition in terms of orthogonal (rather than merely conjugate) vectors. This result allows us to obtain the resolution matrix quickly for the model space from this form. In particular, if we define the diagonal matrix,

$$\mathbf{D}_{\mathbf{G}} = \mathbf{G}_n^{\mathrm{T}} \mathbf{G}_n, \tag{16}$$

we see that

$$\mathbf{M}^{\mathrm{T}}\mathbf{M} \simeq \mathbf{G}_n \mathbf{D}_{\mathbf{G}}{}^{-1} \mathbf{T}_n \mathbf{D}_{\mathbf{G}}{}^{-1} \mathbf{G}_n^{\mathrm{T}}, \tag{17}$$

and therefore, since

$$\mathcal{R}_{\mathrm{model}} \equiv \left(\mathbf{M}^{\mathrm{T}}\mathbf{M}\right)^{\dagger} \mathbf{M}^{\mathrm{T}}\mathbf{M} = \mathbf{M}^{\mathrm{T}}\mathbf{M} \left(\mathbf{M}^{\mathrm{T}}\mathbf{M}\right)^{\dagger}, \tag{18}$$

we find easily that

$$\mathcal{R}_{\mathrm{model}} = \mathbf{G}_n \mathbf{D}_{\mathbf{G}}{}^{-1} \mathbf{G}_n^{\mathrm{T}} = \sum_{i=1}^{n} \frac{\mathbf{g}_i \mathbf{g}_i^{\mathrm{T}}}{(\mathbf{g}_i, \mathbf{g}_i)}. \tag{19}$$

Data resolution estimate

The data resolution is known to be related to the operator,

$$\mathcal{R}_{\mathrm{data}} = \mathbf{M} \left(\mathbf{M}^{\mathrm{T}}\mathbf{M}\right)^{\dagger} \mathbf{M}^{\mathrm{T}}. \tag{20}$$

Substituting (6) for the pseudo-inverse and then defining

$$\mathbf{q}_i \equiv \mathbf{M}\mathbf{p}_i, \tag{21}$$

we find that the resolution operator for the data space is

$$\mathcal{R}_{\mathrm{data}} = \sum_{i=1}^{n} \frac{\mathbf{q}_i \mathbf{q}_i^{\mathrm{T}}}{(\mathbf{q}_i, \mathbf{q}_i)}, \tag{22}$$

a form completely analogous to that in (19).

APPLICATION TO KIRCHHOFF IMAGING

When attempting to image complex subsurfaces with Kirchhoff methods, many difficulties may arise. In particular, amplitude behaviour of the imaged reflectors could be caused by any of several totally different physical phenomenon. Fading and disappearance of a reflector may have several causes, including a real change in reflectivity, an error in the velocity model or an illumination problem. All of these provide

valuable information, but it is important to know which one is causing the effect. By estimating the model resolution it is possible to identify areas of low illumination.

To test the resolution matrix estimation, we inverted for a single output offset (225 m) made from three data offsets (200, 225 and 250 m). The Kirchhoff operator was a simple 2D modelling operator and its adjoint using second-order, first-arrival eikonal traveltimes. For the synthetic case a smoothed version of the correct velocity model was used. The real data example uses a smoothed version of the sequential migration aided reflection tomography (SMART) (Jacobs *et al.* 1992; Ehinger and Lailly 1995) velocity model provided by Elf Aquitaine.

Synthetic test

We began our experiments on the synthetic Elf North Sea data set. This data set was created by Elf, IFP and CGG and inspired by the real data set recorded in the North Sea (block L7d). The synthetic data correspond to a regular 2D acquisition pattern with a source spacing of 50 m, receiver spacing of 25 m and offsets covering a range from 190 to 3340 m.

Figure 1 shows the result of conjugate gradient estimation after 20 iterations. The deepest reflector seems to disappear as it passes under the edge of the salt body. This behaviour is known to be caused by poor illumination.

Figures 2–5 show the estimated resolution for the synthetic data set with an increasing number of iterations. After only five iterations, there is high resolution along the major reflectors (black indicates high resolution, white indicates low resolution). Note that the area of poor illumination has low resolution. As the number of iterations increases, the areas of low reflectivity between the major reflectors become better resolved. This tells us that the conjugate gradient algorithm is expending most of its effort at low iterations resolving model components around the major reflectors. It moves on to the areas of lower reflectivity only at higher iterations.

This behaviour is not surprising, since most of the energy in the model space is found around the major reflectors so that is what will be minimized first.

Real data test

After experimenting with the synthetic data set, we conducted the same experiments on real data (Figs 6 and 7). The data set was acquired by Elf Aquitaine in the North Sea. For the testing purposes, we used one line from the regularized 3D data set with the acquisition parameters similar to those of the synthetic data. Note that the *x*-axis in the real data set is reversed from that in the synthetic so that the salt structure tilts to the left rather than to the right. Figures 8–11 show the results of increasing the iterations for estimating the

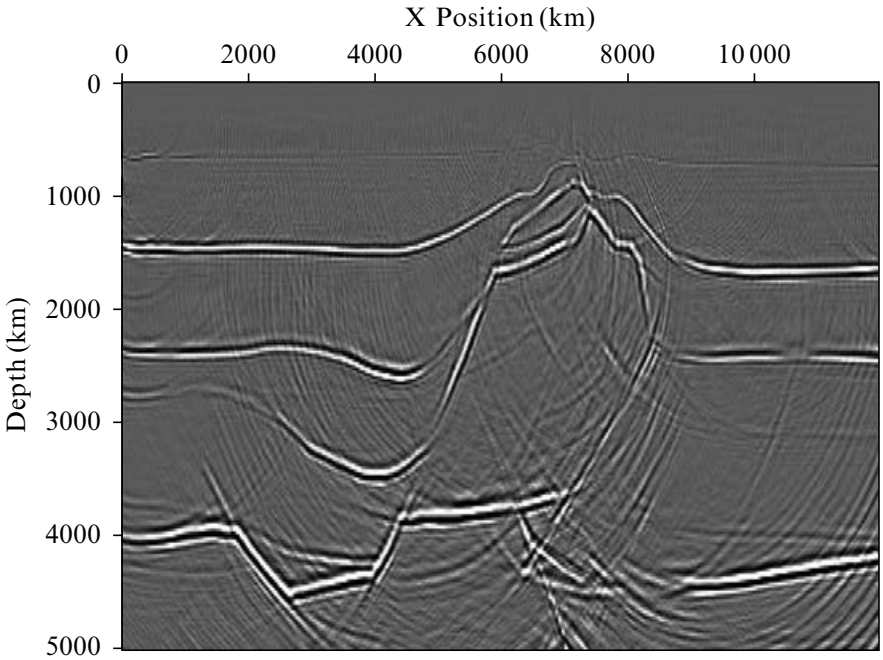

Figure 1 Inversion result on synthetic data after 20 conjugate-gradient iterations.

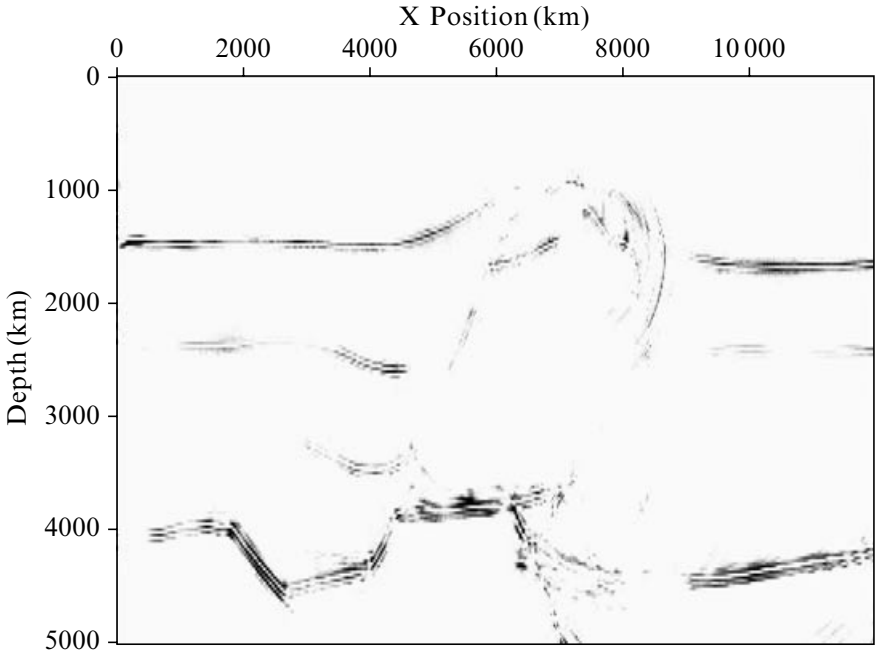

Figure 2 Resolution using the conjugate-gradient method after five iterations. Dark indicates higher resolution.

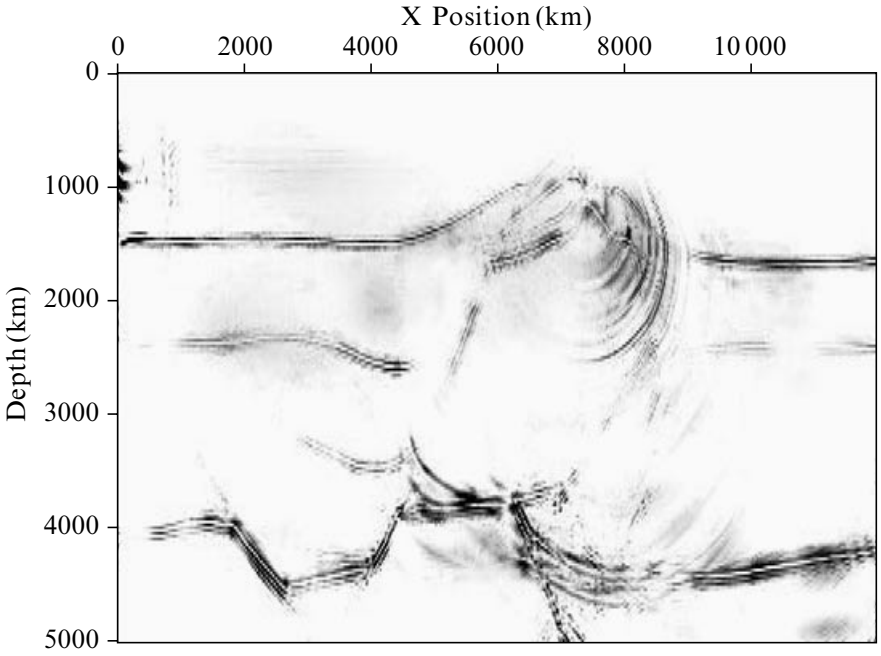

Figure 3 As Fig. 2 after 10 iterations.

resolution. Once again, there are Kirchhoff-type artefacts in all of the figures. Note that we again see resolution energy beginning around the major reflectors, spreading to areas of lower reflectivity at higher iterations.

We can see corresponding changes in our image. After five iterations the image shows strong energy along the primaries reflectors, but it is generally low frequency in Fig. 6. After 20 iterations (Fig. 7), we have an image with more noise, but

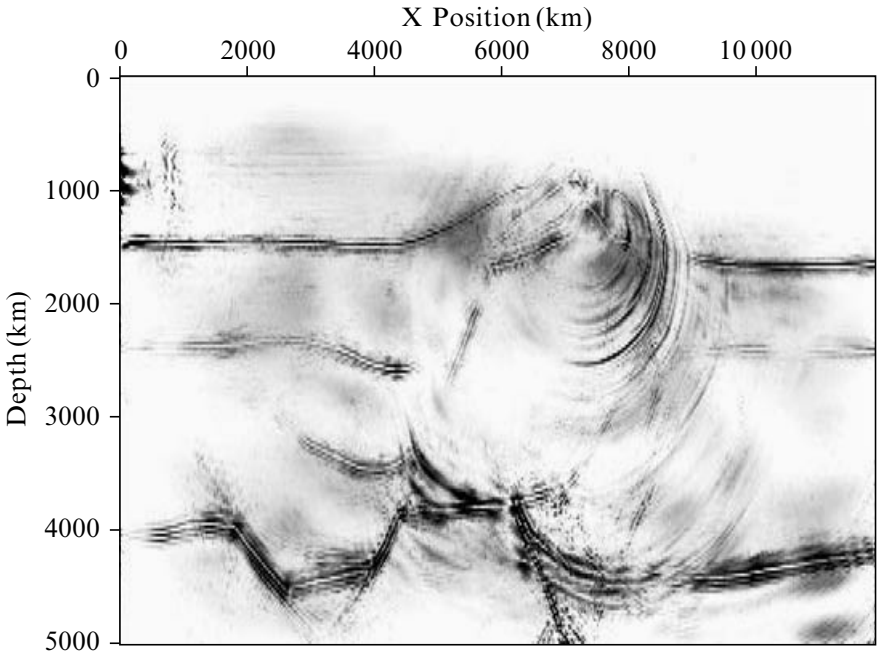

Figure 4 As Fig. 2 after 15 iterations.

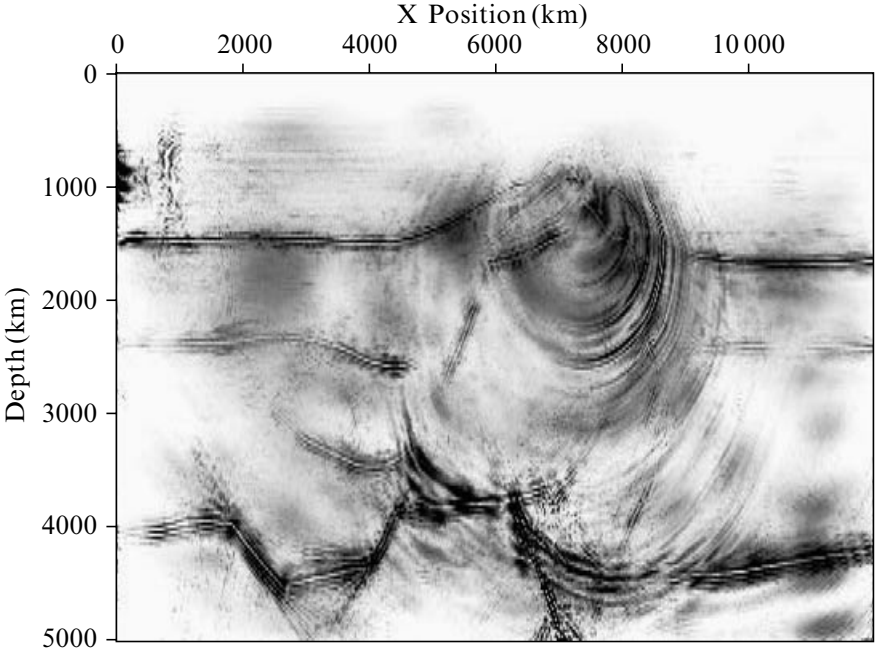

Figure 5 As Fig. 2 after 20 iterations.

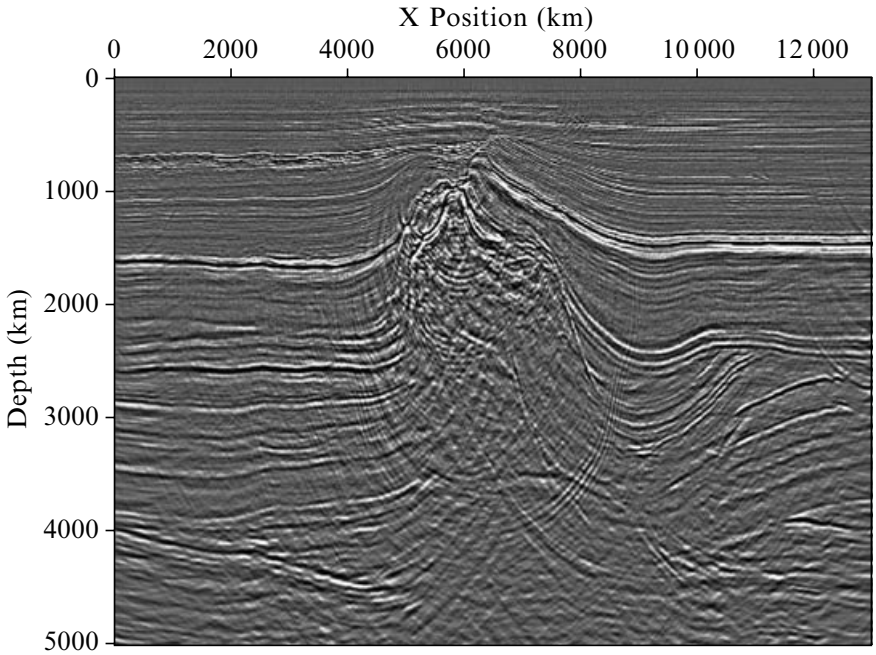

Figure 6 Inversion result on real data after five conjugate-gradient iterations.

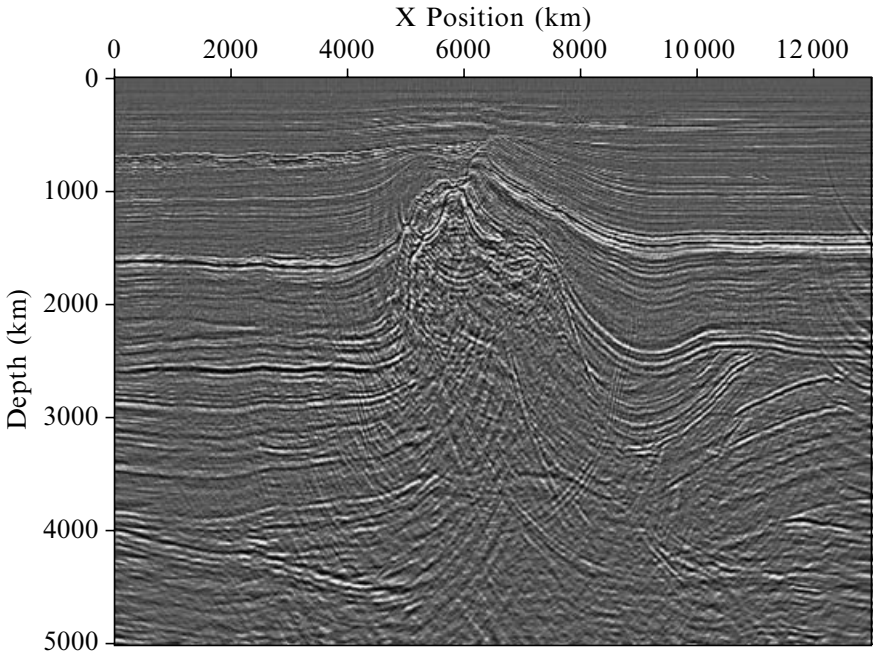

Figure 7 Inversion result on real data after 20 conjugate-gradient iterations.

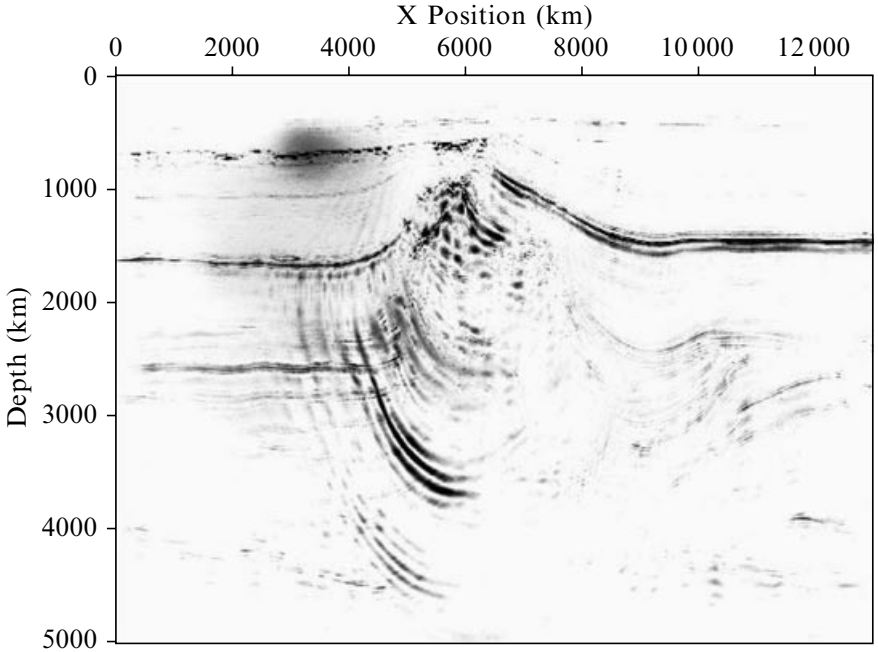

Figure 8 Resolution using the conjugate-gradient method after five iterations of the real data. Dark indicates higher resolution.

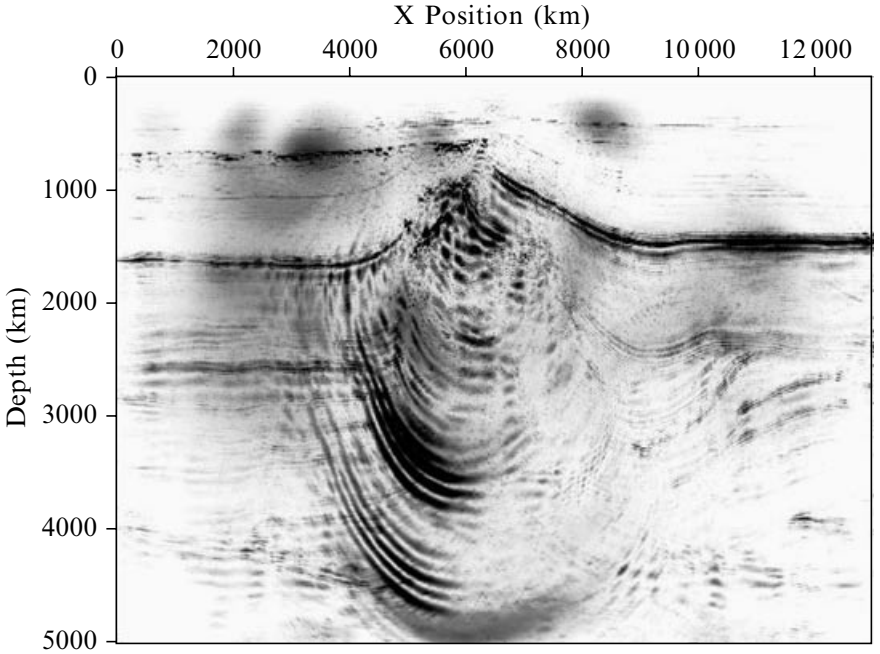

Figure 9 As Fig. 8 after 10 iterations.

also a significantly higher frequency image. The later iterations resolved smaller eigenvalues of the model, which corresponded to higher-frequency, lower-amplitude portions of the model space.

CONCLUSIONS

We have developed a practical approach to resolution estimation using iterative methods of least-squares optimization.

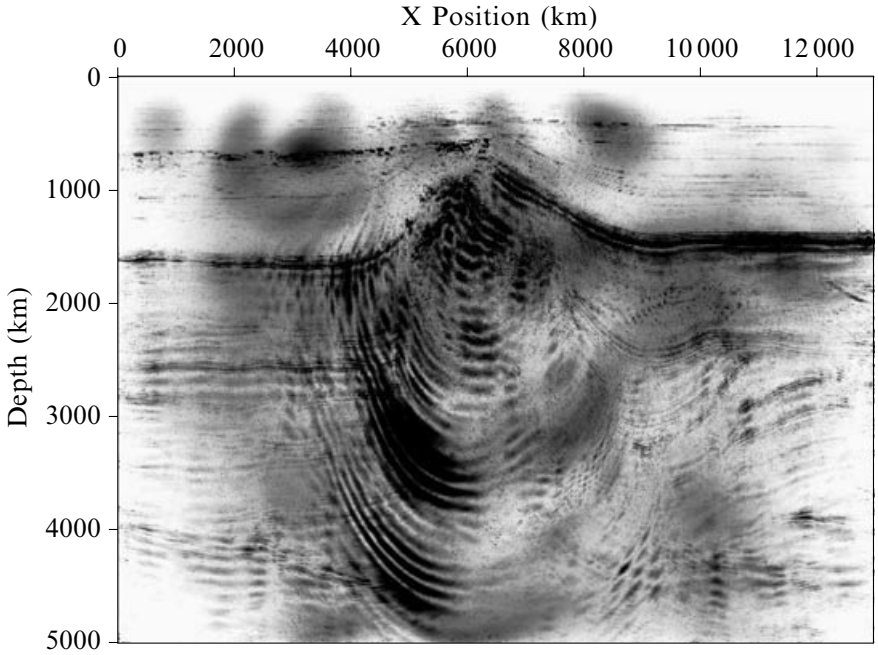

Figure 10 As Fig. 8 after 15 iterations.

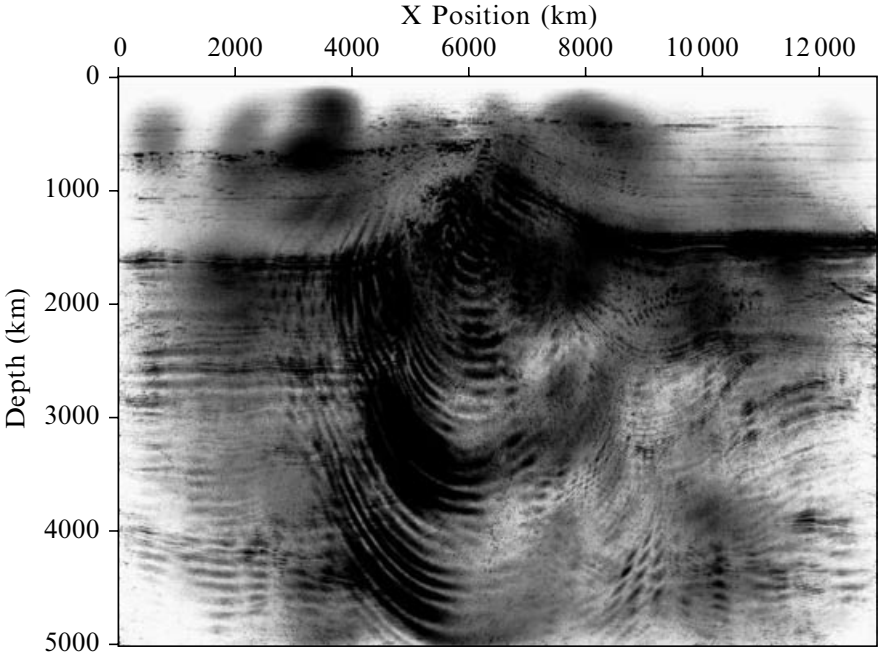

Figure 11 As Fig. 8 after 20 iterations.

Analysing popular iterative methods such as the method of conjugate gradients, we have reformulated the pseudo-inverse, model resolution and data resolution operators in terms of effective iterative estimates.

Iterative estimation of resolution supplies useful information when performing Kirchhoff imaging. Areas of low illumination are easily recognizable. In addition, the iterative nature of the algorithm provides useful information on

what portion of the image is resolvable at early iterations of least-squares migration.

ACKNOWLEDGEMENTS

We thank Elf Aquitaine for providing the data set used in this study. The financial support for our research was provided by the sponsors of the Stanford Exploration Project.

The work of J.G.B. was performed under the auspices of the US Department of Energy by the University of California Lawrence Livermore National Laboratory under contract No. W-7405-ENG-48 and supported specifically by the Geosciences Research Program of the DOE Office of Energy Research within the Office of Basic Energy Sciences, Division of Engineering and Geosciences.

REFERENCES

Bear G., Lu R., Lu C.P., Watson I. and Willen D. 1999. The construction of subsurface illumination and amplitude maps via ray tracing. 69th SEG meeting, Houston, USA, Expanded Abstracts, 1532–1535.

Berryman J.G. 2000a. Analysis of approximate inverses in tomography I. Resolution analysis of common inverses. *Optimization and Engineering* **1**, 87–115.

Berryman J.G. 2000b. Analysis of approximate inverses in tomography II. Iterative inverses. *Optimization and Engineering* **1**, 437–473.

Chavent G. and Plessix R.E. 1999. An optimal true-amplitude least-squares prestack depth-migration operator. *Geophysics* **64**, 508–515.

Claerbout J.F. 1992. *Earth Soundings Analysis: Processing versus Inversion.* Blackwell Scientific Publications.

Cole S. and Karrenbach M. 1992. *Least-squares Kirchhoff Migration.* Report No. 95, pp. 101–110. Stanford Exploration Project. Available at http://sepwww.stanford.edu/public/docs/sep75/toc_html.

Duquet B. and Marfurt K.J. 1999. Filtering coherent noise during prestack depth migration. *Geophysics* **64**, 1054–1066.

Duquet B., Marfurt K.J. and Dellinger J.A. 1998. Efficient estimates of subsurface illumination for Kirchhoff prestack depth migration. 68th SEG meeting, New Orleans, USA, Expanded Abstracts, 1116–1119.

Ehinger A. and Lailly P. 1995. Velocity model determination by the smart method, part 1: Theory. 65th SEG meeting, Houston, USA, Expanded Abstracts, 739–742.

Fletcher R. and Reeves C.M. 1964. Function minimization by conjugate gradients. *Computer Journal* **7**, 149–154.

Hestenes M.R. and Stiefel E. 1952. Methods of conjugate gradients for solving linear systems. *Journal Res. Nat. Bur. Stand.* **49**, 409–436.

Jackson D.D. 1972. Interpretation of inaccurate, insufficient, and inconsistent data. *Geophysical Journal of the Royal Astronomical Society* **28**, 97–109.

Jacobs J.A.C., Delprat-Jannaud F., Ehinger A. and Lailly P. 1992. Sequential migration aided reflection tomography: a tool for imaging complex structures. 62nd SEG meeting, New Orleans, USA, Expanded Abstracts, 1054–1057.

Kleinman R.E., van den and Berg P.M. 1991. Iterative methods for solving integral equations. *Radio Science* **26**, 175–181.

Lanczos C. 1950. An iteration method for the solution of the eigenvalue problem of linear differential and integral operators. *Journal Res. Nat. Bur. Stand.* **45**, 255–282.

Laurain R. and Vinje V. 2001. Prestack depth migration and illumination maps. 71st SEG meeting, San Antonio, USA, Expanded Abstracts, 929–932.

Minkoff S.E. 1996. A computationally feasible approximate resolution matrix for seismic inverse problems. *Geophysical Journal International* **126**, 345–359.

Muerdter D. and Ratcliff D. 2000. Subsalt illumination determined by ray-trace modeling: a catalog of seismic amplitude distortions caused by various salt shapes. 70th SEG meeting, Calgary, Canada, Expanded Abstracts, 733–736.

Nemeth T. 1999. *Imaging and filtering by least-squares migration.* PhD thesis, University of Utah.

Nemeth T., Wu C. and Schuster G.T. 1999. Least-squares migration of incomplete reflection data. *Geophysics* **64**, 208–221.

Paige C.C. and Saunders M.A. 1982. LSQR: An algorithm for sparse linear equations and sparse least squares. *ACM Transactions in Mathematical Software* **8**, 43–71.

Tarantola A. 1987. *Inverse Problem Theory: Methods for Data Fitting and Model Parameter Estimation.* Elsevier Science Publishing Co.

Wang Y.H. and Pratt R.G. 1997. Sensitivities of seismic traveltimes and amplitudes in reflection tomography. *Geophysical Journal International* **131**, 618–642.

Yao Z.S., Roberts R.G. and Tryggvason A. 1999. Calculating resolution and covariance matrices for seismic tomography with the LSQR method. *Geophysical Journal International* **138**, 886–894.

APPENDIX

Review of conjugate directions and conjugate gradients

This Appendix contains a review of the most common methods of iterative least-squares optimization. It explains the background for constructing the iterative resolution estimates.

One useful way to proceed from the iterative equations (2) and (3) is to choose the optimization parameter α_n so that the residual vector is decreased and preferably minimized at each step of the iteration scheme. Using the standard inner Aproduct notation $(\cdot,\cdot)$ and considering

$$||\mathbf{r}_n||^2 = ||\mathbf{r}_{n-1}||^2 - 2\alpha_n(\mathbf{r}_{n-1}, \mathbf{M}\mathbf{p}_{n-1}) + \alpha_n^2||\mathbf{M}\mathbf{p}_{n-1}||^2, \qquad \text{(A1)}$$

we find easily that the optimum choice of α_n using this criterion is

$$\alpha_n = \frac{(\mathbf{r}_{n-1}, \mathbf{M}\mathbf{p}_{n-1})}{||\mathbf{M}\mathbf{p}_{n-1}||^2}. \tag{A2}$$

This formula has the significance that, whenever the residual $\mathbf{r}_{n-1}$ has a component along the direction $\mathbf{M}\mathbf{p}_{n-1}$, α_n is chosen to scale $\mathbf{M}\mathbf{p}_{n-1}$ so that this component exactly cancels and therefore removes the contribution to $\mathbf{r}_n$ made by $\mathbf{M}\mathbf{p}_{n-1}$. This result implies therefore that, if $(\mathbf{r}_{n-1}, \mathbf{M}\mathbf{p}_{n-1}) \neq 0$, then with this choice of α_n we have

$$(\mathbf{r}_n, \mathbf{M}\mathbf{p}_{n-1}) = (\mathbf{M}^{\mathrm{T}}\mathbf{r}_n, \mathbf{p}_{n-1}) = 0. \tag{A3}$$

We used the adjoint property of the inner product in (A3) to show that $\mathbf{p}_{n-1}$ is orthogonal to the gradient vector $\mathbf{g}_n \equiv \mathbf{M}^{\mathrm{T}}\mathbf{r}_n$, so-called because it is the gradient obtained by taking the derivative with respect to $\mathbf{x}_n^{\mathrm{T}}$ of the squared residual error functional associated with (1).

Thus, at each step of this iterative sequence a vector proportional to some vector $\mathbf{p}_n$ is added to the solution, while a vector proportional to $\mathbf{M}\mathbf{p}_n$ is subtracted from the residual. According to formulae (A1) and (A2), the squared norm of the residual decreases at each iteration as

$$||\mathbf{r}_n||^2 = ||\mathbf{r}_{n-1}||^2 - \frac{(\mathbf{r}_{n-1}, \mathbf{M}\mathbf{p}_{n-1})^2}{||\mathbf{M}\mathbf{p}_{n-1}||^2}\ . \tag{A4}$$

The sequence of directions will be most efficient if the vectors used in decimating the residual are orthogonal, i.e. if

$$(\mathbf{M}\mathbf{p}_n, \mathbf{M}\mathbf{p}_j) = 0 \quad \text{for} \quad j = 1, 2, \ldots, n-1. \tag{A5}$$

In this case, as follows by induction from formula (A3), the residual vector is also orthogonal to all those vectors:

$$(\mathbf{r}_n, \mathbf{M}\mathbf{p}_j) = 0 \quad \text{for} \quad j = 1, 2, \ldots, n-1. \tag{A6}$$

Using again the adjoint relation for the inner product, we find that

$$(\mathbf{M}^{\mathrm{T}}\mathbf{r}_n, \mathbf{p}_j) = (\mathbf{g}_n, \mathbf{p}_j) = 0 \quad \text{for} \quad j = 1, 2, \ldots, n-1, \tag{A7}$$

and

$$(\mathbf{p}_n, \mathbf{M}^{\mathrm{T}}\mathbf{M}\mathbf{p}_j) = 0 \quad \text{for} \quad j = 1, 2, \ldots, n-1, \tag{A8}$$

which is a statement of conjugacy for the vectors $\mathbf{p}_n$. Conjugacy is just a generalization of orthogonality in which the vectors are orthogonal relative to the non-standard inner product $(\cdot, \mathbf{A}\cdot)$ – with $\mathbf{A}$ being a symmetric, positive semi-definite matrix (operator) – instead of the standard inner product given by $(\cdot,\cdot)$ with $\mathbf{A}$ replaced by the identity.

We conclude that conjugacy is a desirable property of the set of direction vectors $\mathbf{p}_n$, so our next necessary step in order to obtain a definite iterative process is to construct a convenient sequence of vectors that have this property. One set of model vectors that will be available in this iteration sequence is the set of gradient vectors themselves, where $\mathbf{g}_n = \mathbf{M}^{\mathrm{T}}\mathbf{r}_n$. We show next why this set plays an important role in constructing the desired sequence.

Conjugate directions

To construct a set of directions $\mathbf{p}_n$ which satisfy the conjugacy criterion (A5), we can start from an arbitrary set of model-space vectors $\mathbf{c}_n$ and apply an orthogonalization process to their projections in the data space. An iterative orthogonalization is defined by recursion

$$\mathbf{p}_n = \mathbf{c}_n - \sum_{j=1}^{n-1} \beta_n^{(j)} \mathbf{p}_j\ , \tag{A9}$$

where the following choice of the scalar coefficients $\beta_n^{(j)}$ ensures condition (A5):

$$\beta_n^{(j)} = \frac{(\mathbf{M}\mathbf{c}_n, \mathbf{M}\mathbf{p}_j)}{||\mathbf{M}\mathbf{p}_j||^2}. \tag{A10}$$

Because the residual vector $\mathbf{r}_n$ is orthogonal to all the previous steps in the data space (equation (A6)), the coefficient α_n simplifies to

$$\alpha_n = \frac{(\mathbf{r}_{n-1}, \mathbf{M}\mathbf{c}_{n-1})}{||\mathbf{M}\mathbf{p}_{n-1}||^2}. \tag{A11}$$

Formulae (A9)–(A11) define the method of conjugate directions also known as the preconditioned Krylov subspace method (Kleinman and van den Berg 1991) and under several other names.

A particular choice of the initial directions $\mathbf{c}_n = \mathbf{g}_n = \mathbf{M}^{\mathrm{T}}\mathbf{r}_n$ transforms the method of conjugate directions into the method of conjugate gradients and introduces remarkable simplifications.

Conjugate gradients

First, we notice that the scaling coefficient α_n simplifies with the choice $\mathbf{c}_n = \mathbf{g}_n$ to the form

$$\alpha_n = \frac{(\mathbf{r}_{n-1}, \mathbf{M}\mathbf{g}_{n-1})}{||\mathbf{M}\mathbf{p}_{n-1}||^2} = \frac{(\mathbf{M}^{\mathrm{T}}r_{n-1}, \mathbf{g}_{n-1})}{||\mathbf{M}\mathbf{p}_{n-1}||^2} = \frac{||\mathbf{g}_{n-1}||^2}{||\mathbf{M}\mathbf{p}_{n-1}||^2}, \tag{A12}$$

and the residual decrease (A4) becomes

$$||\mathbf{r}_n||^2 = ||\mathbf{r}_{n-1}||^2 - \frac{||\mathbf{g}_{n-1}||^2}{||\mathbf{M}\mathbf{p}_{n-1}||^2}. \tag{A13}$$

According to (A13), the residual norm is guaranteed to decrease monotonically at each iteration as long as the gradient is different from zero.

Second, applying (A7), we notice the equality

$$(\mathbf{g}_n, \mathbf{g}_j) = (\mathbf{g}_n, \mathbf{p}_j) + \sum_{i=1}^{j-1} \beta_j^{(i)} (\mathbf{g}_n, \mathbf{p}_i) = 0 \quad \text{for} \quad j = 1, 2, \ldots, n-1, \tag{A14}$$

which is precisely equivalent to the conjugacy of the residual vectors. The conjugacy condition can be stated as

$$(\mathbf{M}^{\mathrm{T}}\mathbf{r}_n, \mathbf{M}^{\mathrm{T}}\mathbf{r}_j) = (\mathbf{r}_n, \mathbf{M}\mathbf{M}^{\mathrm{T}}\mathbf{r}_j) = 0 \quad \text{for} \quad j = 1, 2, \ldots, n-1. \tag{A15}$$

Equation (A14) states that the residuals from successive conjugate-gradient iterations form an orthogonal basis in the space of the model $\mathbf{x}$. This fact ensures that the global minimum in an n-dimensional space can be found, in precise arithmetic, in exactly n iterations. We can see that the validity of this remarkable fact is based entirely upon the orthogonality condition (A14).

With $\mathbf{c}_n = \mathbf{g}_n$, we can rewrite (A10) in the form,

$$\beta_n^{(j)} = \frac{(\mathbf{M}\mathbf{g}_n, \mathbf{M}\mathbf{p}_j)}{||\mathbf{M}\mathbf{p}_j||^2} = \frac{(\mathbf{M}\mathbf{g}_n, \mathbf{r}_{j+1} - \mathbf{r}_j)}{\alpha_{j+1}\,||\mathbf{M}\mathbf{p}_j||^2} = \frac{(\mathbf{g}_n, \mathbf{g}_{j+1} - \mathbf{g}_j)}{\alpha_{j+1}\,||\mathbf{M}\mathbf{p}_j||^2}. \tag{A16}$$

It follows immediately from (A16) and the orthogonality condition that

$$\beta_n^{(j)} = 0 \quad \text{for} \quad j = 1, 2, \ldots, n-2\,, \tag{A17}$$

and

$$\beta_n^{(n-1)} = \frac{||\mathbf{g}_n||^2}{\alpha_n\,||\mathbf{M}\mathbf{p}_{n-1}||^2} = \frac{||\mathbf{g}_n||^2}{||\mathbf{g}_{n-1}||^2}. \tag{A18}$$

The latter equality follows from (A12). Thus, the next direction of the conjugate-gradient iteration is completely defined by a linear combination of the current gradient and the previous direction:

$$\mathbf{p}_n = \mathbf{g}_n - \beta_n^{(n-1)}\mathbf{p}_{n-1}. \tag{A19}$$

Equations (A12), (A18) and (A19) provide a complete definition of the classic conjugate-gradient algorithm (Hestenes and Stiefel 1952; Fletcher and Reeves 1964).

Summarizing our derivation, we conclude that the success of the conjugate-direction method is supported by the orthogonality condition (A5). The success of the conjugate-gradient method requires, in addition, the conjugacy condition (A15), which can be expressed in the model space as the orthogonality of the gradients (A14).