Stanford Exploration Project, Report SERGEY, November 9, 2000, pages 53–??

Iterative methods of optimization with application to crosswell tomography

James G. Berryman and Sergey Fomel[1]

ABSTRACT

We review the theory of iterative optimization, revealing the common origin of different optimization methods and reformulating the pseudoinverse, model resolution, and data resolution operators in terms of effective iterative estimates. Examples from crosswell tomography illustrate the theory and suggest efficient methods of its implementation.

INTRODUCTION

Iterative methods of linear unconstrained optimization, such as the method of conjugate gradients (Hestenes and Stiefel, 1952), Lanczos (1950), LSQR (Paige and Saunders, 1982), GMRES (Saad and Schultz, 1986), and some others, have proven to be important tools in geophysical applications. For a given forward modeling operator, predicting the existing data from an unknown model, iterative methods 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 equal to the number of unknown model parameters. In large-scale problems, typical 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 reasonable number of iterations.

When *solving* the inverse problem is replaced by *estimating* 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 Methods and formulations, designed for the complete solutions, are no longer applicable in the case of iterative estimates.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 pseudoinverse operator, model resolution, and data resolution, conventionally associated with SVD decomposition, which becomes infeasible in many large-scale problems.

In this paper, we review different methods of iterative optimization, primarily the method of conjugate directions and the LSQR method. We prove 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. Whenever possible, the orthogonality conditions should be enforced in numerical implementations as a warranty of stable iterative behavior.

[1]**email:** berryman@sep.stanford.edu, sergey@sep.stanford.edu

We show how to define the effective pseudoinverse operator, model and data resolution for iterative methods. Since the exact solution is not available, those definitions apply to effective iterative estimates of the corresponding operators, which were strictly defined in the inversion theory.

Finally, we illustrate the theory with simple examples from crosswell traveltime tomography.

CONJUGATE DIRECTIONS AND CONJUGATE GRADIENTS

We will use a notation consistent with earlier work of Berryman (1994) on crosswell seismic tomography in which the linear inversion problem to be solved takes the form

$$\mathbf{Ms} = \mathbf{t}, \tag{1}$$

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

In crosswell tomography example, $\mathbf{s}^T = (s_1, s_2, \ldots, s_n)$ is an n-vector of wave slownesses associated in either two- or three-dimensions with cells of constant slowness, $\mathbf{M}$ is a matrix of ray-path lengths such that M_{ij} is the length of the i-th ray path through the j-th cell, and $\mathbf{t}^T = (t_1, t_2, \ldots, t_m)$ is an m-vector of the traveltimes associated with the ray paths between specified and numbered pairs of sources and receivers. The assumption that the ray-path matrix $\mathbf{M}$ is known corresponds to assuming that the full inverse problem is being solved in an iterative fashion — in which case the ray-path matrix in question is just the one in use in the latest iteration. We generally assume in addition that the problem is overdetermined so that $m > n$, *i.e.*, the number of data exceed 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

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

where $\mathbf{s}_{n-1}$ is the preceding estimate of $\mathbf{s}$, $\mathbf{s}_n$ is the new estimate of $\mathbf{s}$, $\mathbf{p}_{n-1}$ 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 $\mathbf{r}_n \equiv \mathbf{t} - \mathbf{Ms}_n$, we find the general relation that

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

One useful way to proceed 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 product notation $(\cdot, \cdot)$ and considering

$$||\mathbf{r}_n||^2 = ||\mathbf{r}_{n-1}||^2 - 2\alpha_n(\mathbf{r}_{n-1}, \mathbf{Mp}_{n-1}) + \alpha_n^2 ||\mathbf{Mp}_{n-1}||^2, \tag{4}$$

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

$$\alpha_n = \frac{(\mathbf{r}_{n-1}, \mathbf{M}\mathbf{p}_{n-1})}{||\mathbf{M}\mathbf{p}_{n-1}||^2}. \tag{5}$$

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

$$(\mathbf{r}_n, \mathbf{M}\mathbf{p}_{n-1}) = (\mathbf{M}^T\mathbf{r}_n, \mathbf{p}_{n-1}) = 0. \tag{6}$$

We used the adjoint property of the inner product in (6) to show that $\mathbf{p}_{n-1}$ is orthogonal to the gradient vector $\mathbf{g}_n \equiv \mathbf{M}^T\mathbf{r}_n$, so-called because it is the gradient obtained by taking the derivative with respect to $\mathbf{s}_n^T$ 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 formulas (4) and (5), the squared norm of the residual decreases at each iteration as

$$||\mathbf{r}_n||^2 = ||\mathbf{r}_{n-1}||^2 - \frac{(\mathbf{r}_{n-1}, \mathbf{M}\mathbf{p}_{n-1})^2}{||\mathbf{M}\mathbf{p}_{n-1}||^2} \tag{7}$$

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

$$(\mathbf{M}\mathbf{p}_n, \mathbf{M}\mathbf{p}_j) = 0 \qquad \text{for} \qquad j = 1, 2, \ldots, n-1. \tag{8}$$

In this case, as follows by induction from formula (6), the residual vector is also orthogonal to all those vectors:

$$(\mathbf{r}_n, \mathbf{M}\mathbf{p}_j) = 0 \qquad \text{for} \qquad j = 1, 2, \ldots, n-1. \tag{9}$$

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

$$(\mathbf{M}^T\mathbf{r}_n, \mathbf{p}_j) = (\mathbf{g}_n, \mathbf{p}_j) = 0 \qquad \text{for} \qquad j = 1, 2, \ldots, n-1. \tag{10}$$

and

$$(\mathbf{p}_n, \mathbf{M}^T\mathbf{M}\mathbf{p}_j) = 0 \qquad \text{for} \qquad j = 1, 2, \ldots, n-1, \tag{11}$$

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 the nonstandard inner product $(\cdot, \mathbf{A}\cdot)$ – with $\mathbf{A}$ being a symmetric, positive semidefinite 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 $\mathbf{g}_n = \mathbf{M}^T\mathbf{r}_n$. We show next why this set plays an important role in constructing the desired sequence.

Why not orthogonal residuals?

Since the direction vectors have not yet been specified, we still have some degree of freedom that we may use to help choose an optimum method. We might suppose that it could be possible to choose the residuals themselves in a way so that they are orthogonal from one step to the next. But, we soon find this is wrong for, if we were to try this, we would find that the condition

$$(\mathbf{r}_n, \mathbf{r}_{n-1}) = 0 \tag{12}$$

implies that

$$\alpha_n = \frac{(\mathbf{r}_{n-1}, \mathbf{r}_{n-1})}{(\mathbf{r}_{n-1}, \mathbf{M}\mathbf{p}_{n-1})} = \frac{(\mathbf{r}_{n-1}, \mathbf{M}\mathbf{p}_{n-1})}{||\mathbf{M}\mathbf{p}_{n-1}||^2}. \tag{13}$$

But, this condition (13) is actually *impossible* because the Cauchy-Schwartz inequality for vectors states that

$$(\mathbf{r}_{n-1}, \mathbf{M}\mathbf{p}_{n-1})^2 \leq ||\mathbf{r}_{n-1}||^2 \, ||\mathbf{M}\mathbf{p}_{n-1}||^2, \tag{14}$$

where the equality in (14) occurs only when the two vectors are proportional which will virtually never be true for such an iteration scheme.

Since orthogonality of the residuals is impossible, the next best condition we might try to impose is conjugacy of the residual vectors. This condition can be stated as

$$(\mathbf{M}^T\mathbf{r}_n, \mathbf{M}^T\mathbf{r}_j) = (\mathbf{r}_n, \mathbf{M}\mathbf{M}^T\mathbf{r}_j) = 0 \qquad \text{for} \qquad j = 1, 2, \ldots, n-1. \tag{15}$$

Thus, if this condition is met, the residuals will be orthogonal relative to the inner product $(\cdot, \mathbf{M}\mathbf{M}^T\cdot)$. The set of conditions (15) is not in conflict with (5) so we assume it holds and construct the directions $\mathbf{p}_n$ that satisfy both (11) and (15).

Conjugate directions

To construct a set of directions $\mathbf{p}_n$ which satisfy the conjugacy criterion (8), 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

$$\mathbf{p}_n = \mathbf{c}_n - \sum_{j=1}^{n-1} \beta_n^{(j)} \mathbf{p}_j \ , \tag{16}$$

where the following choice of the scalar coefficients $\beta_n^{(j)}$ assures condition (8):

$$\beta_n^{(j)} = \frac{(\mathbf{M}\mathbf{c}_n, \mathbf{M}\mathbf{p}_j)}{||\mathbf{M}\mathbf{p}_j||^2}. \tag{17}$$

According to the fact that the residual vector $\mathbf{r}_n$ is orthogonal to all the previous steps in the data space (equation (9)), the coefficient α_n simplifies to

$$\alpha_n = \frac{(\mathbf{r}_{n-1}, \mathbf{M}\mathbf{c}_{n-1})}{||\mathbf{M}\mathbf{p}_{n-1}||^2}. \tag{18}$$

Formulas (16–18) define the method of conjugate directions (Fomel, 1996) 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 $\mathbf{c}_n = \mathbf{g}_n = \mathbf{M}^T\mathbf{r}_n$ 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

$$\alpha_n = \frac{(\mathbf{r}_{n-1}, \mathbf{M}\mathbf{g}_{n-1})}{||\mathbf{M}\mathbf{p}_{n-1}||^2} = \frac{(\mathbf{M}^T r_{n-1}, \mathbf{g}_{n-1})}{||\mathbf{M}\mathbf{p}_{n-1}||^2} = \frac{||\mathbf{g}_{n-1}||^2}{||\mathbf{M}\mathbf{p}_{n-1}||^2}, \tag{19}$$

and the residual decrease (7) becomes

$$||\mathbf{r}_n||^2 = ||\mathbf{r}_{n-1}||^2 - \frac{||\mathbf{g}_{n-1}||^2}{||\mathbf{M}\mathbf{p}_{n-1}||^2} \tag{20}$$

According to formula (20), the residual norm is guaranteed to decrease monotonically at each iteration as long as the gradient is different from zero.

Second, applying formula (10), we notice the equality

$$(\mathbf{g}_n, \mathbf{g}_j) = (\mathbf{g}_n, \mathbf{p}_j) + \sum_{i=1}^{j-1} \beta_j^{(i)} (\mathbf{g}_n, \mathbf{p}_i) = 0 \qquad \text{for} \qquad j = 1, 2, \ldots, n-1\,, \tag{21}$$

which is precisely equivalent to the conjugacy of the residual vectors, suggested earlier by (15). Equation (21) states that the residuals from successive conjugate-gradient iterations form an orthogonal basis in the space of the model $\mathbf{s}$. This fact assures 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 (21).

With $\mathbf{c}_n = \mathbf{g}_n$, we can rewrite equation (17) in the form

$$\beta_n^{(j)} = \frac{(\mathbf{M}\mathbf{g}_n, \mathbf{M}\mathbf{p}_j)}{||\mathbf{M}\mathbf{p}_j||^2} = \frac{(\mathbf{M}\mathbf{g}_n, \mathbf{r}_{j+1} - \mathbf{r}_j)}{\alpha_{j+1}\,||\mathbf{M}\mathbf{p}_j||^2} = \frac{(\mathbf{g}_n, \mathbf{g}_{j+1} - \mathbf{g}_j)}{\alpha_{j+1}\,||\mathbf{M}\mathbf{p}_j||^2}. \tag{22}$$

It follows immediately from formula (22) and the orthogonality condition that

$$\beta_n^{(j)} = 0 \qquad \text{for} \qquad j = 1, 2, \ldots, n-2\,, \tag{23}$$

and

$$\beta_n^{(n-1)} = \frac{||\mathbf{g}_n||^2}{\alpha_n\,||\mathbf{M}\mathbf{p}_{n-1}||^2} = \frac{||\mathbf{g}_n||^2}{||\mathbf{g}_{n-1}||^2}. \tag{24}$$

The latter equality follows from formula (19). Thus, the next direction of the conjugate-gradient iteration is completely defined by a linear combination of the current gradient and the previous direction:

$$\mathbf{p}_n = \mathbf{g}_n - \beta_n^{(n-1)}\mathbf{p}_{n-1}. \tag{25}$$

Equations (19), (24), and (25) 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 (8). The success of the conjugate-gradient method requires, in addition, the conjugacy condition (15), which can be expressed in the model-space as the orthogonality of the gradients (21).

The next section shows how the orthogonal sets of vectors in the data and model spaces translate into the effective resolution operators.

RESOLUTION OPERATORS FOR BOTH MODEL AND DATA

Pseudoinverse estimate

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

$$\mathbf{s}_n = \sum_{i=1}^{n-1} \alpha_{i+1}\mathbf{p}_i \ , \tag{26}$$

where we assume for simplicity that $\mathbf{s}_1 = 0$. Then substituting (5) — or more directly the first ratio in (19) — for the α_i's shows that the n-th iterate is given explicitly by

$$\mathbf{s}_n = \sum_{j=1}^{n-1} \frac{\mathbf{p}_j\mathbf{p}_j^T}{(\mathbf{p}_j, \mathbf{M}^T\mathbf{M}\mathbf{p}_j)}\mathbf{g}_j = \sum_{j=1}^{n-1} \frac{\mathbf{p}_j\mathbf{p}_j^T}{(\mathbf{p}_j, \mathbf{M}^T\mathbf{M}\mathbf{p}_j)}\mathbf{M}^T\mathbf{t} \tag{27}$$

for this scheme. The resulting approximate inverse operator is therefore

$$\left(\mathbf{M}^T\mathbf{M}\right)^\dagger \simeq \sum_{j=1}^{n-1} \frac{\mathbf{p}_j\mathbf{p}_j^T}{(\mathbf{p}_j, \mathbf{M}^T\mathbf{M}\mathbf{p}_j)}, \tag{28}$$

which form we now want to study. We use the dagger notation to indicate that the expression in (28) is approximating a pseudoinverse, because it may happen that the normal matrix is singular in which case the standard inverse does not exist.

First, note that although (28) 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 n-th iteration to be

$$\mathbf{P}_n = \begin{pmatrix} \mathbf{p}_1 & \mathbf{p}_2 & \cdots & \mathbf{p}_n \end{pmatrix}, \tag{29}$$

then the approximate inverse operator can be rewritten as

$$\left(\mathbf{M}^T\mathbf{M}\right)^\dagger \simeq \mathbf{P}_n {\mathbf{D}_\mathbf{P}}^{-1} \mathbf{P}_n^T, \tag{30}$$

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

$$\mathbf{D_P} \equiv \mathbf{P}_n^T \mathbf{M}^T \mathbf{M} \mathbf{P}_n, \tag{31}$$

because of the conjugacy of the $\mathbf{p}$'s composing $\mathbf{P}_n$. Now equation (25) shows that

$$\mathbf{P}_n \mathbf{B}_n = \mathbf{G}_n, \tag{32}$$

where

$$\mathbf{G}_n = \begin{pmatrix} \mathbf{g}_1 & \mathbf{g}_2 & \cdots & \mathbf{g}_n \end{pmatrix}, \tag{33}$$

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

$$\mathbf{B}_n = \begin{bmatrix} 1 & \beta_2^{(1)} & 0 & \cdots & 0 \\ 0 & 1 & \beta_3^{(2)} & \cdots & \cdots \\ 0 & 0 & \cdots & \cdots & 0 \\ \cdots & \cdots & \cdots & 1 & \beta_n^{(n-1)} \\ 0 & \cdots & 0 & 0 & 1 \end{bmatrix} \tag{34}$$

Multiplying (32) on the right by the inverse of $\mathbf{B}_n$ and then substituting into (30), we find that

$$\left(\mathbf{M}^T\mathbf{M}\right)^\dagger \simeq \mathbf{G}_n \mathbf{B}_n^{-1} {\mathbf{D}_\mathbf{P}}^{-1} \left(\mathbf{B}_n^T\right)^{-1} \mathbf{G}_n^T. \tag{35}$$

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

$$\left(\mathbf{M}^T\mathbf{M}\right)^\dagger \simeq \mathbf{G}_n \mathbf{T}_n^{-1} \mathbf{G}_n^T, \tag{36}$$

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

$$\mathbf{T}_n = \mathbf{B}_n^T \mathbf{D_P} \mathbf{B}_n. \tag{37}$$

This result highlights the similarities between the CG method and that of other iterative methods such as Lanczos (1950) and LSQR (Paige and Saunders, 1982), also producing tridiagonal respresentations of the matrix to be inverted.

Model resolution estimate

Although the tridiagonal form found in (37) is interesting in its own right, the more important result contained in (36) 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

$$\mathbf{D_G} = \mathbf{G}_n^T \mathbf{G}_n, \tag{38}$$

we see that

$$\mathbf{M}^T\mathbf{M} \simeq \mathbf{G}_n {\mathbf{D_G}}^{-1} \mathbf{T}_n {\mathbf{D_G}}^{-1} \mathbf{G}_n^T, \tag{39}$$

and therefore, since

$$\mathcal{R}_{model} \equiv \left(\mathbf{M}^T\mathbf{M}\right)^\dagger \mathbf{M}^T\mathbf{M} = \mathbf{M}^T\mathbf{M}\left(\mathbf{M}^T\mathbf{M}\right)^\dagger, \tag{40}$$

we find easily that

$$\mathcal{R}_{model} = \mathbf{G}_n {\mathbf{D_G}}^{-1} \mathbf{G}_n^T = \sum_{i=1}^{n} \frac{\mathbf{g}_i \mathbf{g}_i^T}{(\mathbf{g}_i, \mathbf{g}_i)}. \tag{41}$$

Data resolution estimate

The data resolution is known to be related to the operator

$$\mathcal{R}_{data} = \mathbf{M}\left(\mathbf{M}^T\mathbf{M}\right)^\dagger \mathbf{M}^T. \tag{42}$$

Substituting (28) for the pseudoinverse and then defining

$$\mathbf{q}_i \equiv \mathbf{M}\mathbf{p}_i, \tag{43}$$

we find that the resolution operator for the data space is

$$\mathcal{R}_{data} = \sum_{i=1}^{n} \frac{\mathbf{q}_i \mathbf{q}_i^T}{(\mathbf{q}_i, \mathbf{q}_i)}, \tag{44}$$

a form completely analogous to that in (41).

IMPORTANCE OF MAINTAINING ORTHOGONALITY

As we have been discussing, suppose we want to invert a square, $n \times n$ symmetric matrix $\mathbf{A}$ and we want to do so using some iterative method like Lanczos (1950), LSQR (Paige and

Saunders, 1982), conjugate gradients (Hestenes and Stiefel, 1952), etc. Then the iterative algorithm (especially for Lanczos and LSQR) generally may be expressed in the form

$$\mathbf{A}\mathbf{Z}_k = \mathbf{Z}_k\mathbf{T}_k + N_{k+1}\mathbf{z}_{k+1}\mathbf{e}_k^T \tag{45}$$

for $k \leq r$ where r is the rank of the matrix $\mathbf{A}$ and the matrix $\mathbf{Z}_k$, given by

$$\mathbf{Z}_k = \begin{pmatrix} \mathbf{z}_1 & \mathbf{z}_2 & \dots & \mathbf{z}_k \end{pmatrix}, \tag{46}$$

is composed of the column vectors $\mathbf{z}_i$ for $i = 1,\dots,k$ that have been generated so far in the iterative process. The other terms in (45) are the unit vector $\mathbf{e}_k$, with its single component in the kth position, and a norming constant N_{k+1}. The formula (45) with $\mathbf{T}$ being a tridiagonal matrix is just exactly the iterative scheme of Lanczos.

Now the model resolution matrix $\mathcal{R}_{model}$ for the iterative scheme at the kth iteration is given by

$$\mathcal{R}_k = \mathbf{Z}_k\mathbf{Z}_k^T. \tag{47}$$

However, care must be taken to make sure that the $\mathbf{z}_i$'s are orthogonal as they are expected to be in this scheme. The tridiagonalization process produces a sequence of orthogonal vectors in principle, but in practice the orthogonality may break down after several iterations when the calculations are performed at finite precision. To demonstrate how the lack of orthogonality affects the process consider the following facts. Let the eigenvalues of the matrix $\mathbf{A}$ be ordered so that $\lambda_i^2 \neq 0$ for $i = 1,\dots,r$ and $\lambda_i^2 = 0$ for $i = r+1,\dots,n$. (We write the eigenvalues as squares because $\mathbf{A}$ often takes the form of a normal matrix $\mathbf{A} = \mathbf{M}^T\mathbf{M}$, in which case the eigenvalues of $\mathbf{M}$ are the λ_is.) Then the trace of the matrix is just

$$\mathrm{Tr}\mathbf{A} = \sum_{i=1}^{r} \lambda_i^2. \tag{48}$$

The model resolution $\mathcal{R}_k$ is a projection operator onto a k-dimensional Krylov subspace (the one explored so far by the iterative method) of the n-dimensional space that is both the range and domain of $\mathbf{A}$. Taking the trace of $\mathcal{R}_k\mathbf{A}$ shows that

$$\mathrm{Tr}\left[\mathbf{Z}_k\mathbf{Z}_k^T\mathbf{A}\right] = \mathrm{Tr}\left[\mathbf{Z}_k^T\mathbf{A}\mathbf{Z}_k\right] = \sum_{i=1}^{k} \mathbf{z}_i^T\mathbf{A}\mathbf{z}_i. \tag{49}$$

Let $\mathbf{v}_i$ for $i = 1,\dots,r$ be the normalized eigenvectors associated with the eigenvalues λ_i^2. Then, assuming only that $\mathbf{z}_1$ has no components in the null-space of $\mathbf{A}$, each of the iteration vectors $\mathbf{z}_i$ can be expanded in terms of these eigenvectors with the result that

$$\mathbf{z}_i = \sum_{j=1}^{r} \zeta_{ij}\mathbf{v}_j. \tag{50}$$

Similarly, the eigenvectors can be expanded in terms of the full set of normalized iteration vectors $\mathbf{z}_i$ according to

$$\mathbf{v}_j = \sum_{i=1}^{r} \zeta_{ij}\mathbf{z}_i, \tag{51}$$

where the same set of coefficients ζ_{ij} is used in both expansions and these coefficients must satisfy

$$\sum_{i=1}^{r} \zeta_{ij}^2 = 1 = \sum_{j=1}^{r} \zeta_{ij}^2 \tag{52}$$

in order for both sets of vectors to be normalized to unity.

Substituting (50) into (49), we find easily that

$$\mathrm{Tr}\left[\mathbf{Z}_k \mathbf{Z}_k^T \mathbf{A}\right] = \sum_{j=1}^{r} \lambda_j^2 \left[\sum_{i=1}^{k} \zeta_{ij}^2 \right] \leq \sum_{j=1}^{r} \lambda_j^2 . \tag{53}$$

The inequality follows, since the coefficients ζ_{ij} are all real, and therefore

$$\sum_{j=1}^{k} \zeta_{ij}^2 \leq \sum_{j=1}^{r} \zeta_{ij}^2 \equiv 1. \tag{54}$$

Thus, we can expect that as long as the vectors returned by the iterative scheme are orthogonal, the effective trace of the operator $\mathbf{T}$ must remain smaller than that of the original operator $\mathbf{A}$. If this constraint is violated, then we know that the set of $\mathbf{z}_i$'s generated so far are not orthogonal and furthermore that this lack of orthogonality is having serious negative consequences on our ability to compute the resolution operator correctly.

This fact leads to a test for orthogonality that is very easy to implement when our operator is a matrix with known elements. We will make use of this test in the next section.

CROSSWELL TOMOGRAPHY APPLICATION

Figures 1-4 provide some numerical examples comparing and contrasting the results obtained using standard SVD resolution calculations with the LSQR (Paige and Saunders, 1982) resolution calculations described in an earlier paper (Berryman, 1994). We consider a 4×4 model using strictly crosswell data, so there are 16 source/receiver pairs as well as 16 cells in 2D. Model slowness values are shown in the upper block of each figure, while diagonal resolution values are shown in the lower block. The first two examples (Figures 1 and 2) show results for the actual model used to compute the traveltime data [see Berryman (1990) for a description of the code used to generate both the forward and inverse solutions]. The second two examples (Figures 3 and 4) show results obtained after 15 iterations of the reconstruction code of Berryman (1990). The LSQR resolution examples (Figures 2 and 4) were computed using ten iterations of the LSQR algorithm, so the maximum size of the resolved model vector space has dimension ten. To aid in the comparison, the SVD resolution examples use only the 10 eigenvectors associated with the 10 largest eigenvalues of the ray-path matrix. We find the results are in qualitatively good agreement. Better quantitative agreement is not anticipated because the 10-dimensional vector spaces spanned by these two approximations, although having large regions of overlap, will nevertheless almost always differ to some degree.

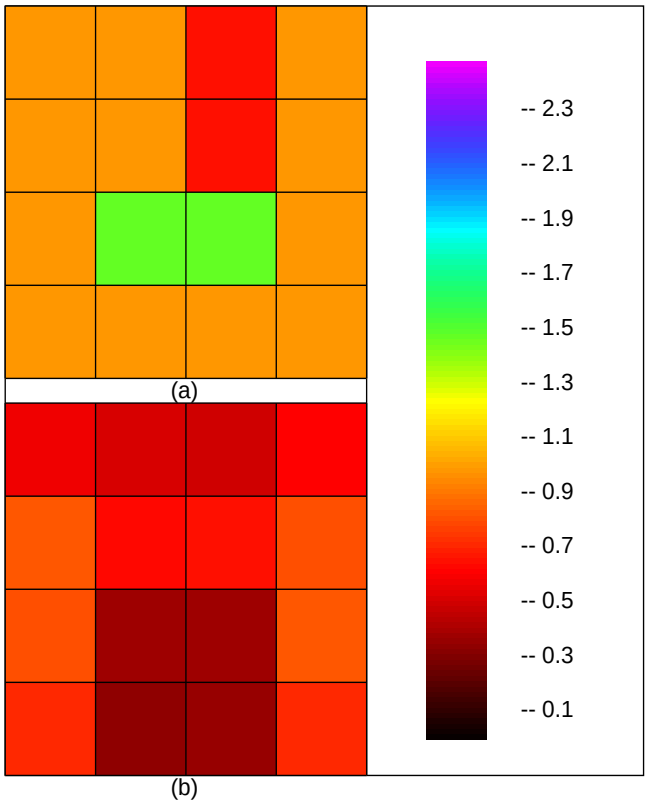

Figure 1: Target model slowness (a) and resolution (b) for truncated SVD using 10 largest eigenvalues. resol-casea [NR]

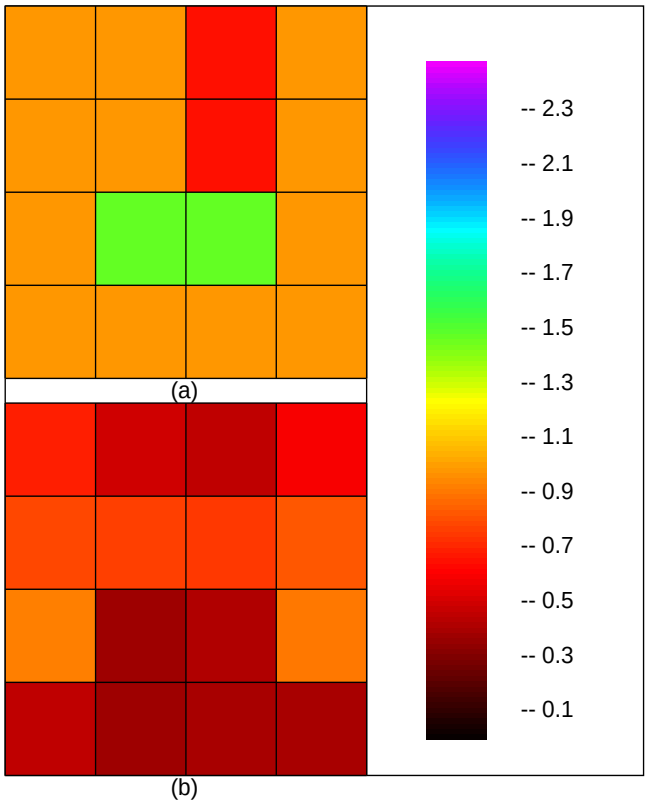

Figure 2: Target model slowness (a) and resolution (b) for LSQR after 10 iteratons. resol-caseb [NR]

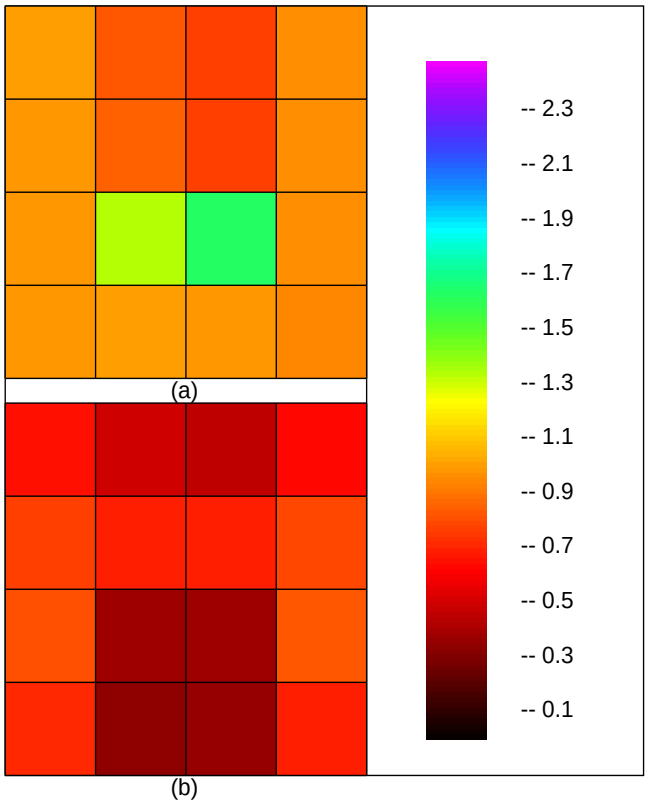

Figure 3: Slowness (a) and resolution (b) for truncated SVD using 10 largest eigenvalues. resol-casec [NR]

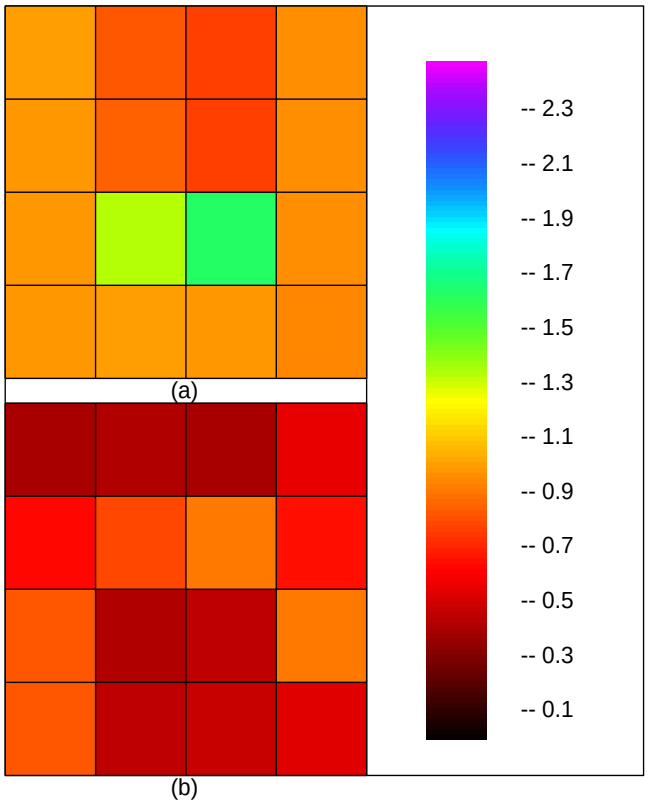

Figure 4: Reconstructed slowness (a) and resolution (b) for LSQR after 10 iterations. resol-cased [NR]

Tests of orthogonalization

For the LSQR algorithm, it is easy to check the trace of the effective operator since it is just the sum of the squares of the elements in the resulting bidiagonal matrix. To test the ideas of the preceding section, we first perform an LSQR inversion while imposing full reorthogonalization (Arnoldi, 1951) at each step. Doing so, we obtain Figure 5, showing that the effective operator trace converges monotonically to the true trace from below as expected.

In contrast, Figure 6 shows that without reorthogonalization the trace of the effective operator has already exceeded the upper bound — thus demonstrating that the eighth direction vector found by LSQR has at least some component parallel to one of the earlier vectors and also along some eigenvector with large eigenvalue. Observing the later behavior of this curve, regions of small slopes may correspond to allowable (orthogonal) direction vectors, but the regions of large slope must be resampling the regions of the vector space with the largest eigenvalues. This observation shows that the rebirth of nonorthogonal vectors does not happen just at the eighth iteration, but is a recurring problem.

These two examples show clearly that full reorthogonalization works very well and that failure to do any reorthogonalization can lead to serious problems with the set of direction vectors generated by such schemes. To make progress, we want to know whether full reorthogonalization is required, or whether some type of partial reothogonalization (which will presumably be cheaper) might be equally or almost as effective as full reorthogonalization. We want to explore the tradeoffs between cost of the partial reorthogonalization and the benefits to be derived from it. Figures 7–11 explore these issues.

Figures 7–9 refer to the same model considered in Figures 1-6. Figure 7 shows that orthogonalizing against the first 1, 2, or 3 vectors improves the results progressively, the more vectors are used for the reorthogonalization. Figure 8 shows that reorthogonalization against the 1, 2, or 3 most recently generated vectors does not work as well as the previous Figure. Reorthogonalizing against both 1, 2, or 3 early and late vectors gives virtually identical results in Figure 9 as those results observed in Figure 7.

Figures 10 and 11 refer to a model of 16×8 cells similar to that considered earlier by Berryman (1990). Figure 10 shows that reorthogonalizing against only the first and last vectors generated in the iteration sequence is quite ineffective at reducing the nonorthogonal vectors generated. Figure 11 shows that reorthogonalizing against the first 35 vectors produces a major improvement, without significant orthogonalization problems out to 90 iterations, instead of less than 35 iterations before problems arise with only one vector reothogonalization.

We conclude that reorthogonalization is effective and the partial reorthogonalization is most effective when the vectors chosen for the reorthogonalization set are those from the early part of the iteration sequence. The reason that these vectors are best to use is presumably because they correspond to directions that have components parallel to directions in the space that are eigenvectors of the operator being inverted with largest eigenvalues. These vectors like to be reborn in this process and reorthogonalization is an effective means of preventing multiply copies of the same dominant vectors from recurring in the iteration sequence.

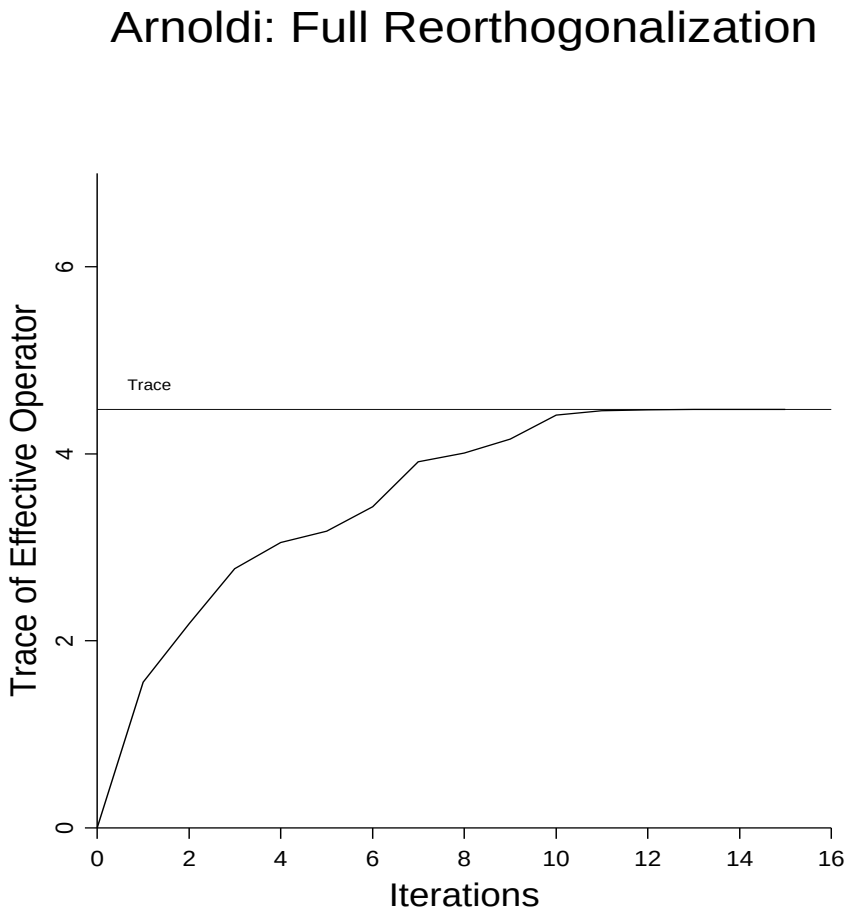

Figure 5: LSQR with full (Arnoldi) reorthogonalization. resol-sumarnoldi [NR]

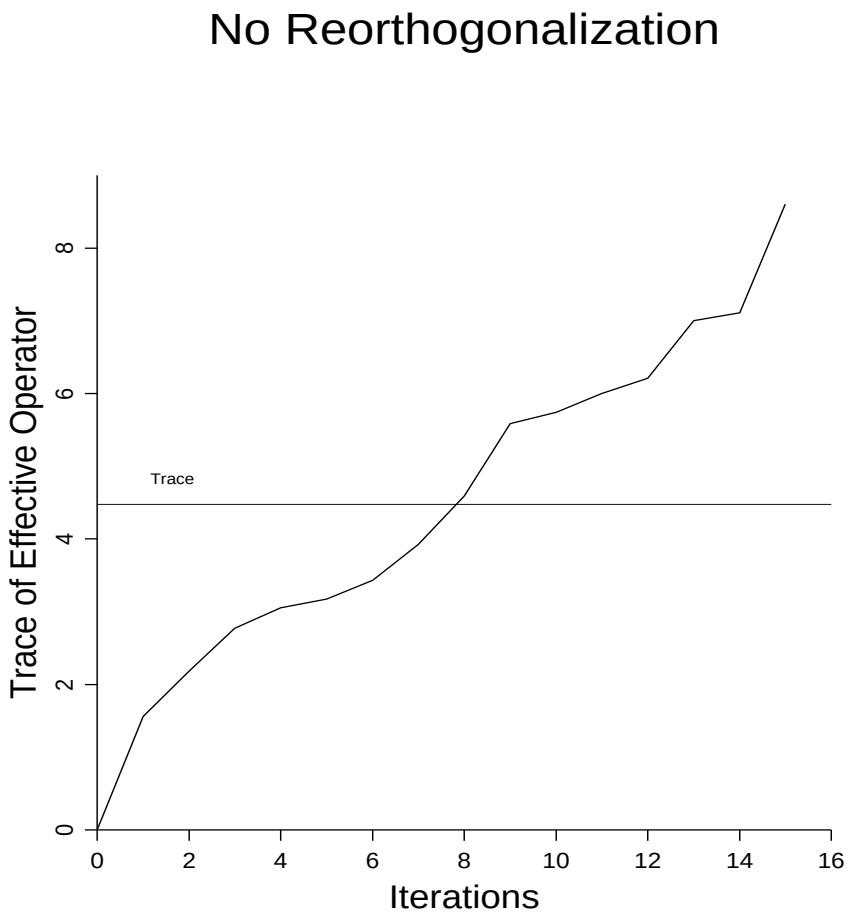

Figure 6: LSQR with no reorthogonalization. resol-sum000 [NR]

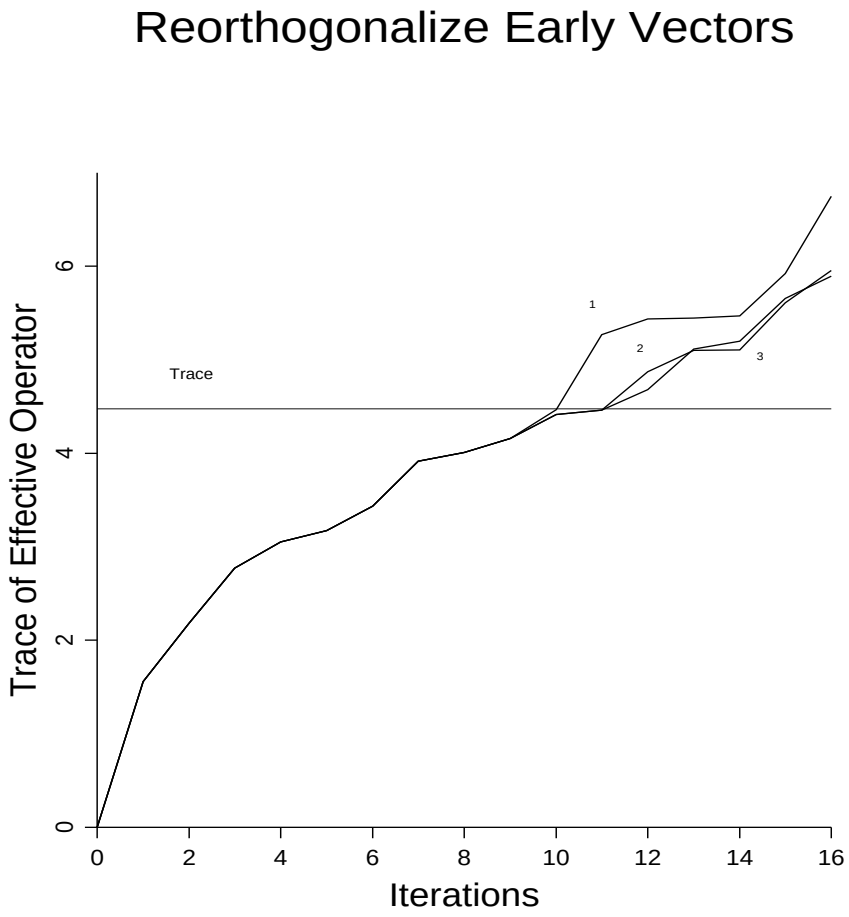

Figure 7: LSQR with partial reorthogonalization – first one, two, or three. resol-sum310 [NR]

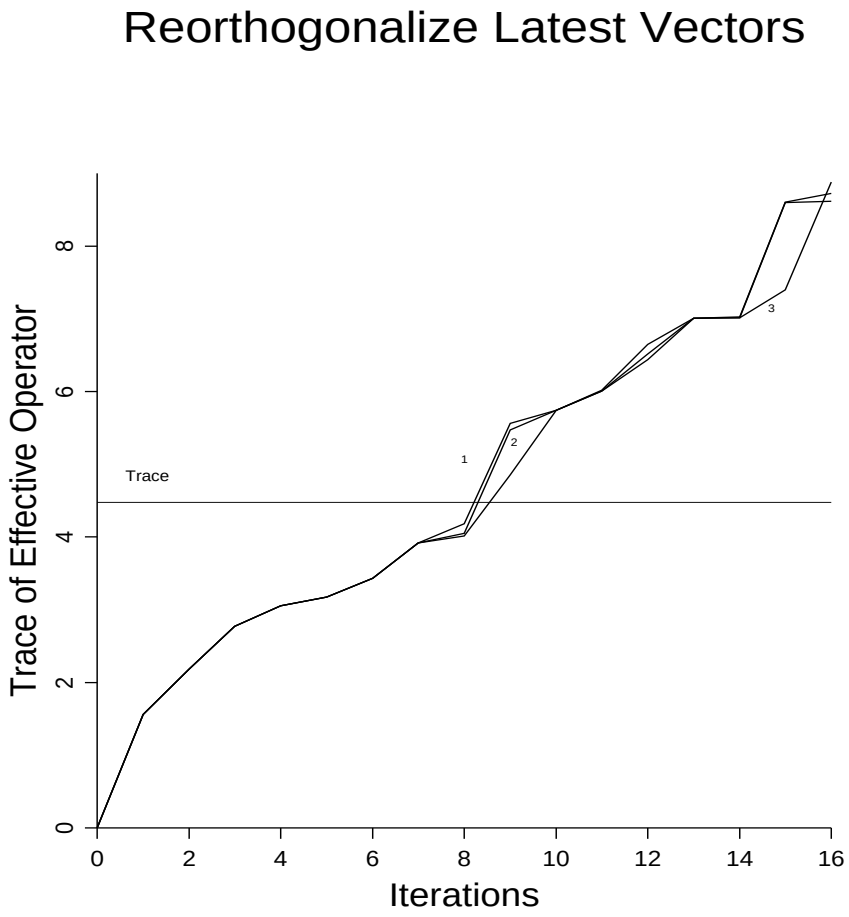

Figure 8: LSQR with partial reorthogonalization – last one, two, or three. resol-sum301 [NR]

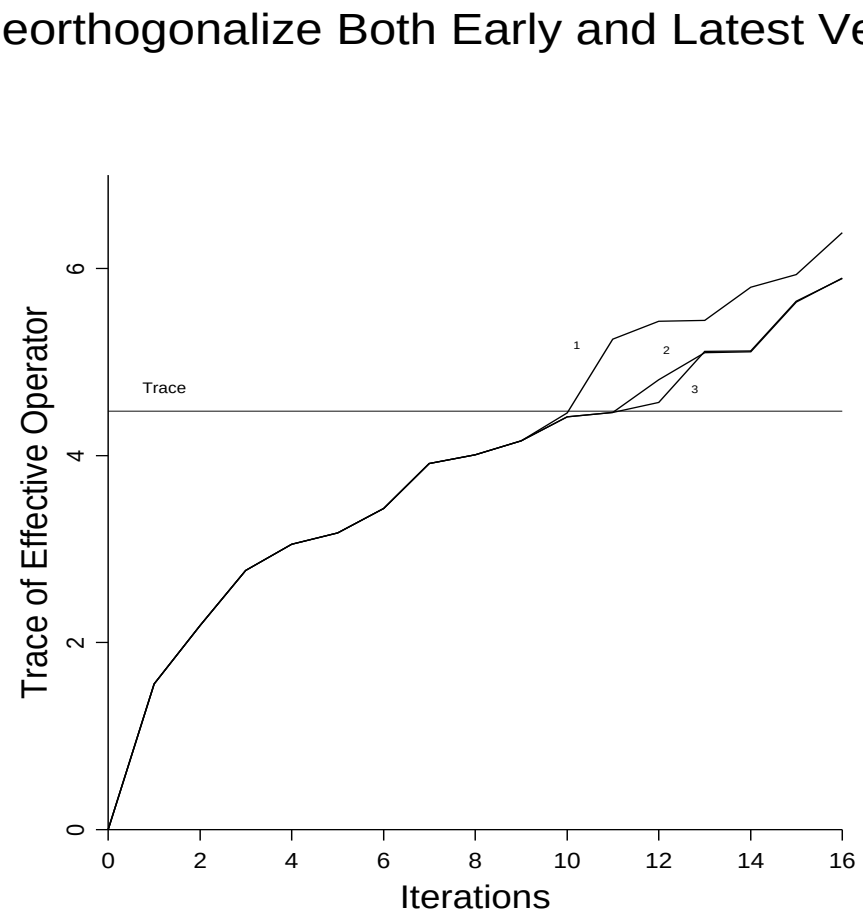

Figure 9: LSQR with partial reorthogonalization – using both first and last one, two, or three. resol-sum311 [NR]

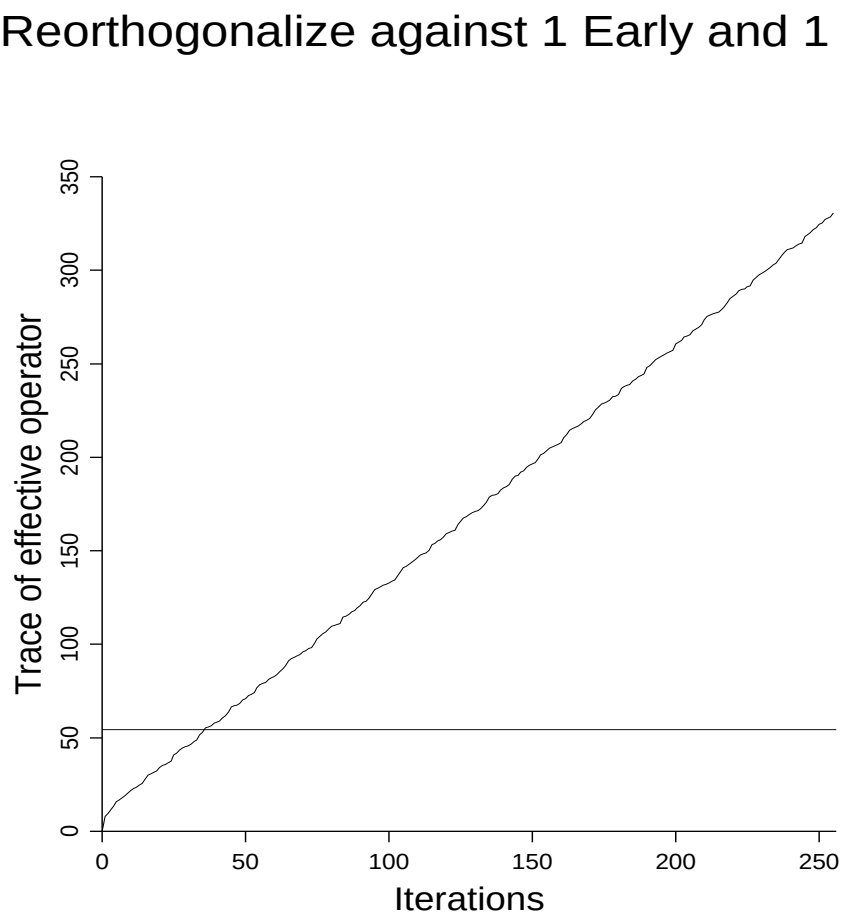

Figure 10: LSQR with partial reorthogonalization – using first and last vector. resol-ex111 [NR]

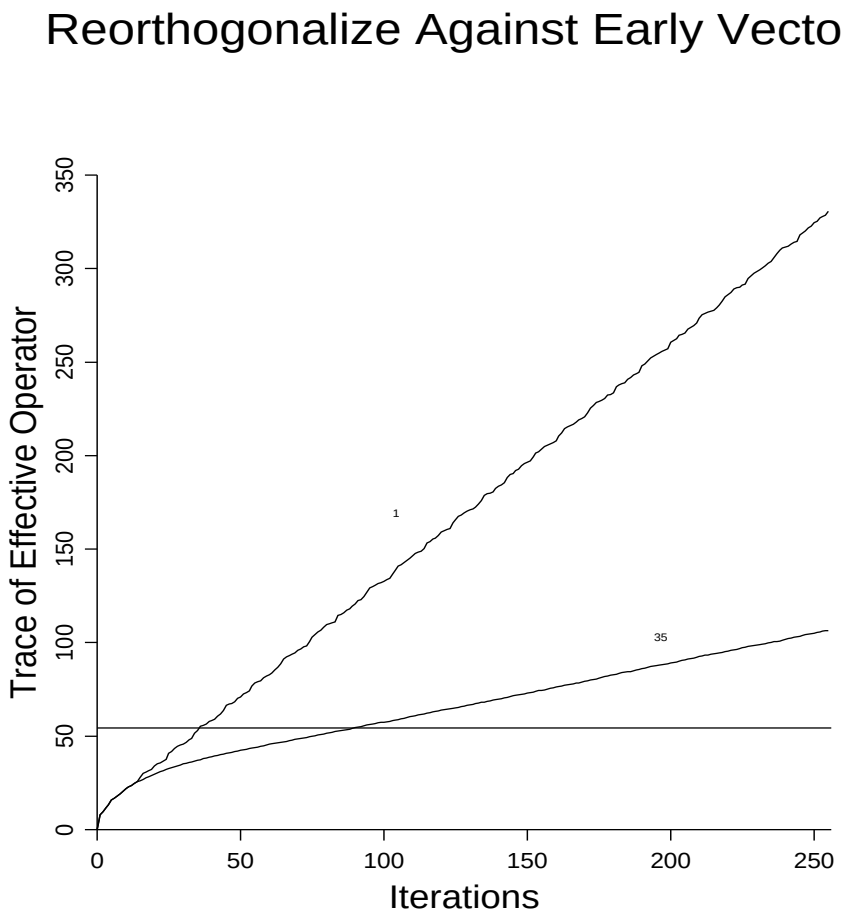

Figure 11: LSQR with partial reorthogonalization – using first or first 35 vectors. resol-ex3510 [NR]

CONCLUSIONS

The methods of computing resolution matrices that have been illustrated here may be easily generalized to a variety of other iterative and approximate inversion methods. We have explored partial reorthogonalization methods for iterative methods and have found that using a subset of the early vectors generated in the iteration sequence is most effective at reducing unwanted occurrences of nonorthogonal vectors in the later parts of the iteration sequence. These early vectors correspond to directions that have components along the eigenvectors with the largest eigenvalues, and these are precisely the vectors we most need to exclude from the later iterations. Such recurrences may not adversely affect the inversion itself, but do make the computation of the resolution matrices (operators) much more complicated than if the orthogonalization is enforced.

REFERENCES

Arnoldi, W. E., 1951, The principle of minimized iterations in the solution of the matrix eigenvalue problem: Quart. Appl. Math., **9**, 17–29.

Berryman, J. G., 1990, Stable iterative reconstruction algorithm for nonlinear traveltime tomography: Inverse Problems, **6**, 21–42.

Berryman, J. G., 1994, Tomographic resolution without singular value decomposition: Mathematical Methods in Geophysical Imaging II, 1–13.

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

Fomel, S., 1996, Least-square inversion with inexact adjoints. Method of conjugate directions: A tutorial: SEP–**92**, 253–265.

Hestenes, M. R., and Stiefel, E., 1952, Methods of conjugate gradients for solving linear systems: J. Res. NBS, **49**, 409–436.

Kleinman, R. E., and van den 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 opertors: J. Res. Nat. Bur. Stand., **45**, 255–282.

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

Saad, Y., and Schultz, M., 1986, GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems: SIAM J. Sci. Statist. Comput., **7**, 856–869.