

# Least-square inversion with inexact adjoints. Method of conjugate directions: A tutorial<sup>a</sup>

<sup>a</sup>Published in SEP Report, 92, 253-365 (1996)

*Sergey Fomel*<sup>1</sup>

## ABSTRACT

This tutorial describes the classic method of conjugate directions: the generalization of the conjugate-gradient method in iterative least-square inversion. I derive the algebraic equations of the conjugate-direction method from general optimization principles. The derivation explains the “magic” properties of conjugate gradients. It also justifies the use of conjugate directions in cases when these properties are distorted either by computational errors or by inexact adjoint operators. The extra cost comes from storing a larger number of previous search directions in the computer memory. A simple program and two examples illustrate the method.

## INTRODUCTION

This paper describes the method of conjugate directions for solving linear operator equations in Hilbert space. This method is usually described in the numerous textbooks on unconstrained optimization as an introduction to the much more popular method of conjugate gradients. See, for example, *Practical optimization* by Gill et al. (1995) and its bibliography. The famous conjugate-gradient solver possesses specific properties, well-known from the original works of Hestenes and Stiefel (1952) and Fletcher and Reeves (1964). For linear operators and exact computations, it guarantees finding the solution after, at most,  $n$  iterative steps, where  $n$  is the number of dimensions in the solution space. The method of conjugate gradients doesn't require explicit computation of the objective function and explicit inversion of the Hessian matrix. This makes it particularly attractive for large-scale inverse problems, such as those of seismic data processing and interpretation. However, it does require explicit computation of the adjoint operator. Claerbout (1992, 2003) shows dozens of successful examples of the conjugate gradient application with numerically precise adjoint operators.

The motivation for this tutorial is to explore the possibility of using different types of preconditioning operators in the place of adjoints in iterative least-square inversion. For some linear or linearized operators, implementing the exact adjoint may pose a

---

<sup>1</sup>**e-mail:** sergey@sep.stanford.edu

difficult problem. For others, one may prefer different preconditioners because of their smoothness (Claerbout, 1995a; Crawley, 1995), simplicity (Kleinman and van den Berg, 1991), or asymptotic properties (Sevink and Herman, 1994). In those cases, we could apply the natural generalization of the conjugate gradient method, which is the method of conjugate directions. The cost difference between those two methods is in the volume of memory storage. In the days when the conjugate gradient method was invented, this difference looked too large to even consider a practical application of conjugate directions. With the evident increase of computer power over the last 30 years, we can afford to do it now.

I derive the main equations used in the conjugate-direction method from very general optimization criteria, with minimum restrictions implied. The textbook algebra is illustrated with a simple program and two simple examples.

## IN SEARCH OF THE MINIMUM

We are looking for the solution of the linear operator equation

$$\mathbf{d} = \mathbf{A} \mathbf{m} , \quad (1)$$

where  $\mathbf{m}$  is the unknown model in the linear model space,  $\mathbf{d}$  stands for the given data, and  $\mathbf{A}$  is the forward modeling operator. The data vector  $\mathbf{d}$  belongs to a Hilbert space with a defined norm and dot product. The solution is constructed by iterative steps in the model space, starting from an initial guess  $\mathbf{m}_0$ . Thus, at the  $n$ -th iteration, the current model  $\mathbf{m}_n$  is found by the recursive relation

$$\mathbf{m}_n = \mathbf{m}_{n-1} + \alpha_n \mathbf{s}_n , \quad (2)$$

where  $\mathbf{s}_n$  denotes the step direction, and  $\alpha_n$  stands for the scaling coefficient. The residual at the  $n$ -th iteration is defined by

$$\mathbf{r}_n = \mathbf{d} - \mathbf{A} \mathbf{m}_n . \quad (3)$$

Substituting (2) into (3) leads to the equation

$$\mathbf{r}_n = \mathbf{r}_{n-1} - \alpha_n \mathbf{A} \mathbf{s}_n . \quad (4)$$

For a given step  $\mathbf{s}_n$ , we can choose  $\alpha_n$  to minimize the squared norm of the residual

$$\|\mathbf{r}_n\|^2 = \|\mathbf{r}_{n-1}\|^2 - 2 \alpha_n (\mathbf{r}_{n-1}, \mathbf{A} \mathbf{s}_n) + \alpha_n^2 \|\mathbf{A} \mathbf{s}_n\|^2 . \quad (5)$$

The parentheses denote the dot product, and  $\|\mathbf{x}\| = \sqrt{(\mathbf{x}, \mathbf{x})}$  denotes the norm of  $x$  in the corresponding Hilbert space. The optimal value of  $\alpha_n$  is easily found from equation (5) to be

$$\alpha_n = \frac{(\mathbf{r}_{n-1}, \mathbf{A} \mathbf{s}_n)}{\|\mathbf{A} \mathbf{s}_n\|^2} . \quad (6)$$

Two important conclusions immediately follow from this fact. First, substituting the value of  $\alpha_n$  from formula (6) into equation (4) and multiplying both sides of this equation by  $\mathbf{r}_n$ , we can conclude that

$$(\mathbf{r}_n, \mathbf{A} \mathbf{s}_n) = 0, \quad (7)$$

which means that the new residual is orthogonal to the corresponding step in the residual space. This situation is schematically shown in Figure 1. Second, substituting formula (6) into (5), we can conclude that the new residual decreases according to

$$\|\mathbf{r}_n\|^2 = \|\mathbf{r}_{n-1}\|^2 - \frac{(\mathbf{r}_{n-1}, \mathbf{A} \mathbf{s}_n)^2}{\|\mathbf{A} \mathbf{s}_n\|^2}, \quad (8)$$

(“Pythagoras’s theorem” ), unless  $\mathbf{r}_{n-1}$  and  $\mathbf{A} \mathbf{s}_n$  are orthogonal. These two conclusions are the basic features of optimization by the method of steepest descent. They will help us define an improved search direction at each iteration.




Figure 1: Geometry of the residual in the data space (a scheme).

## IN SEARCH OF THE DIRECTION

Let’s suppose we have a generator that provides particular search directions at each step. The new direction can be the gradient of the objective function (as in the method of steepest descent), some other operator applied on the residual from the previous step, or, generally speaking, any arbitrary vector in the model space. Let us denote the automatically generated direction by  $\mathbf{c}_n$ . According to formula (8), the residual decreases as a result of choosing this direction by

$$\|\mathbf{r}_{n-1}\|^2 - \|\mathbf{r}_n\|^2 = \frac{(\mathbf{r}_{n-1}, \mathbf{A} \mathbf{c}_n)^2}{\|\mathbf{A} \mathbf{c}_n\|^2}. \quad (9)$$

How can we improve on this result?

## First step of the improvement

Assuming  $n > 1$ , we can add some amount of the previous step  $\mathbf{s}_{n-1}$  to the chosen direction  $\mathbf{c}_n$  to produce a new search direction  $\mathbf{s}_n^{(n-1)}$ , as follows:

$$\mathbf{s}_n^{(n-1)} = \mathbf{c}_n + \beta_n^{(n-1)} \mathbf{s}_{n-1}, \quad (10)$$

where  $\beta_n^{(n-1)}$  is an adjustable scalar coefficient. According to the fundamental orthogonality principle (7),

$$(\mathbf{r}_{n-1}, \mathbf{A} \mathbf{s}_{n-1}) = 0. \quad (11)$$

As follows from equation (11), the numerator on the right-hand side of equation (9) is not affected by the new choice of the search direction:

$$(\mathbf{r}_{n-1}, \mathbf{A} \mathbf{s}_n^{(n-1)})^2 = [(\mathbf{r}_{n-1}, \mathbf{A} \mathbf{c}_n) + \beta_n^{(n-1)} (\mathbf{r}_{n-1}, \mathbf{A} \mathbf{s}_{n-1})]^2 = (\mathbf{r}_{n-1}, \mathbf{A} \mathbf{c}_n)^2. \quad (12)$$

However, we can use transformation (10) to decrease the denominator in (9), thus further decreasing the residual  $\mathbf{r}_n$ . We achieve the minimization of the denominator

$$\|\mathbf{A} \mathbf{s}_n^{(n-1)}\|^2 = \|\mathbf{A} \mathbf{c}_n\|^2 + 2\beta_n^{(n-1)} (\mathbf{A} \mathbf{c}_n, \mathbf{A} \mathbf{s}_{n-1}) + (\beta_n^{(n-1)})^2 \|\mathbf{A} \mathbf{s}_{n-1}\|^2 \quad (13)$$

by choosing the coefficient  $\beta_n^{(n-1)}$  to be

$$\beta_n^{(n-1)} = -\frac{(\mathbf{A} \mathbf{c}_n, \mathbf{A} \mathbf{s}_{n-1})}{\|\mathbf{A} \mathbf{s}_{n-1}\|^2}. \quad (14)$$

Note the analogy between (14) and (6). Analogously to (7), equation (14) is equivalent to the orthogonality condition

$$(\mathbf{A} \mathbf{s}_n^{(n-1)}, \mathbf{A} \mathbf{s}_{n-1}) = 0. \quad (15)$$

Analogously to (8), applying formula (14) is also equivalent to defining the minimized denominator as

$$\|\mathbf{A} \mathbf{s}_n^{(n-1)}\|^2 = \|\mathbf{A} \mathbf{c}_n\|^2 - \frac{(\mathbf{A} \mathbf{c}_n, \mathbf{A} \mathbf{s}_{n-1})^2}{\|\mathbf{A} \mathbf{s}_{n-1}\|^2}. \quad (16)$$

## Second step of the improvement

Now let us assume  $n > 2$  and add some amount of the step from the  $(n-2)$ -th iteration to the search direction, determining the new direction  $\mathbf{s}_n^{(n-2)}$ , as follows:

$$\mathbf{s}_n^{(n-2)} = \mathbf{s}_n^{(n-1)} + \beta_n^{(n-2)} \mathbf{s}_{n-2}. \quad (17)$$

We can deduce that after the second change, the value of numerator in equation (9) is still the same:

$$(\mathbf{r}_{n-1}, \mathbf{A} \mathbf{s}_n^{(n-2)})^2 = [(\mathbf{r}_{n-1}, \mathbf{A} \mathbf{c}_n) + \beta_n^{(n-2)} (\mathbf{r}_{n-1}, \mathbf{A} \mathbf{s}_{n-2})]^2 = (\mathbf{r}_{n-1}, \mathbf{A} \mathbf{c}_n)^2. \quad (18)$$

This remarkable fact occurs as the result of transforming the dot product  $(\mathbf{r}_{n-1}, \mathbf{A} \mathbf{s}_{n-2})$  with the help of equation (4):

$$(\mathbf{r}_{n-1}, \mathbf{A} \mathbf{s}_{n-2}) = (\mathbf{r}_{n-2}, \mathbf{A} \mathbf{s}_{n-2}) - \alpha_{n-1} (\mathbf{A} \mathbf{s}_{n-1}, \mathbf{A} \mathbf{s}_{n-2}) = 0. \quad (19)$$

The first term in (19) is equal to zero according to formula (7); the second term is equal to zero according to formula (15). Thus we have proved the new orthogonality equation

$$(\mathbf{r}_{n-1}, \mathbf{A} \mathbf{s}_{n-2}) = 0, \quad (20)$$

which in turn leads to the numerator invariance (18). The value of the coefficient  $\beta_n^{(n-2)}$  in (17) is defined analogously to (14) as

$$\beta_n^{(n-2)} = -\frac{(\mathbf{A} \mathbf{s}_n^{(n-1)}, \mathbf{A} \mathbf{s}_{n-2})}{\|\mathbf{A} \mathbf{s}_{n-2}\|^2} = -\frac{(\mathbf{A} \mathbf{c}_n, \mathbf{A} \mathbf{s}_{n-2})}{\|\mathbf{A} \mathbf{s}_{n-2}\|^2}, \quad (21)$$

where we have again used equation (15). If  $\mathbf{A} \mathbf{s}_{n-2}$  is not orthogonal to  $\mathbf{A} \mathbf{c}_n$ , the second step of the improvement leads to a further decrease of the denominator in (8) and, consequently, to a further decrease of the residual.

## Induction

Continuing by induction the process of adding a linear combination of the previous steps to the arbitrarily chosen direction  $\mathbf{c}_n$  (known in mathematics as the *Gram-Schmidt orthogonalization process*), we finally arrive at the complete definition of the new step  $\mathbf{s}_n$ , as follows:

$$\mathbf{s}_n = \mathbf{s}_n^{(1)} = \mathbf{c}_n + \sum_{j=1}^{j=n-1} \beta_n^{(j)} \mathbf{s}_j. \quad (22)$$

Here the coefficients  $\beta_n^{(j)}$  are defined by equations

$$\beta_n^{(j)} = -\frac{(\mathbf{A} \mathbf{c}_n, \mathbf{A} \mathbf{s}_j)}{\|\mathbf{A} \mathbf{s}_j\|^2}, \quad (23)$$

which correspond to the orthogonality principles

$$(\mathbf{A} \mathbf{s}_n, \mathbf{A} \mathbf{s}_j) = 0, \quad 1 \leq j \leq n-1 \quad (24)$$

and

$$(\mathbf{r}_n, \mathbf{A} \mathbf{s}_j) = 0, \quad 1 \leq j \leq n. \quad (25)$$

It is these orthogonality properties that allowed us to optimize the search parameters one at a time instead of solving the  $n$ -dimensional system of optimization equations for  $\alpha_n$  and  $\beta_n^{(j)}$ .

## ALGORITHM

The results of the preceding sections define the method of conjugate directions to consist of the following algorithmic steps:

1. Choose initial model  $\mathbf{m}_0$  and compute the residual  $r_0 = \mathbf{d} - \mathbf{A} \mathbf{m}_0$ .
2. At  $n$ -th iteration, choose the initial search direction  $\mathbf{c}_n$ .
3. If  $n$  is greater than 1, optimize the search direction by adding a linear combination of the previous directions, according to equations (22) and (23), and compute the modified step direction  $\mathbf{s}_n$ .
4. Find the step length  $\alpha_n$  according to equation (6). The orthogonality principles (24) and (7) can simplify this equation to the form

$$\alpha_n = \frac{(\mathbf{r}_{n-1}, \mathbf{A} \mathbf{c}_n)}{\|\mathbf{A} \mathbf{s}_n\|^2}. \quad (26)$$

5. Update the model  $\mathbf{m}_n$  and the residual  $\mathbf{r}_n$  according to equations (2) and (4).
6. Repeat iterations until the residual decreases to the required accuracy or as long as it is practical.

At each of the subsequent steps, the residual is guaranteed not to increase according to equation (8). Furthermore, optimizing the search direction guarantees that the convergence rate doesn't decrease in comparison with (9). The only assumption we have to make to arrive at this conclusion is that the operator  $\mathbf{A}$  is linear. However, without additional assumptions, we cannot guarantee global convergence of the algorithm to the least-square solution of equation (1) in a finite number of steps.

## WHAT ARE ADJOINTS FOR? THE METHOD OF CONJUGATE GRADIENTS

The adjoint operator  $\mathbf{A}^T$  projects the data space back to the model space and is defined by the dot product test

$$(\mathbf{d}, \mathbf{A} \mathbf{m}) \equiv (\mathbf{A}^T \mathbf{d}, \mathbf{m}) \quad (27)$$

for any  $\mathbf{m}$  and  $\mathbf{d}$ . The method of conjugate gradients is a particular case of the method of conjugate directions, where the initial search direction  $\mathbf{c}_n$  is

$$\mathbf{c}_n = \mathbf{A}^T \mathbf{r}_{n-1}. \quad (28)$$

This direction is often called the *gradient*, because it corresponds to the local gradient of the squared residual norm with respect to the current model  $\mathbf{m}_{n-1}$ . Aligning the initial search direction along the gradient leads to the following remarkable simplifications in the method of conjugate directions.

## Orthogonality of the gradients

The orthogonality principle (25) transforms according to the dot-product test (27) to the form

$$(\mathbf{r}_{n-1}, \mathbf{A} \mathbf{s}_j) = (\mathbf{A}^T \mathbf{r}_{n-1}, \mathbf{s}_j) = (\mathbf{c}_n, \mathbf{s}_j) = 0, \quad 1 \leq j \leq n-1. \quad (29)$$

Forming the dot product  $(\mathbf{c}_n, \mathbf{c}_j)$  and applying formula (22), we can see that

$$(\mathbf{c}_n, \mathbf{c}_j) = \left( \mathbf{c}_n, \mathbf{s}_j - \sum_{i=1}^{i=j-1} \beta_n^{(i)} \mathbf{s}_i \right) = (\mathbf{c}_n, \mathbf{s}_j) - \sum_{i=1}^{i=j-1} \beta_n^{(i)} (\mathbf{c}_n, \mathbf{s}_i) = 0, \quad 1 \leq j \leq n-1. \quad (30)$$

Equation (30) proves the orthogonality of the gradient directions from different iterations. Since the gradients are orthogonal, after  $n$  iterations they form a basis in the  $n$ -dimensional space. In other words, if the model space has  $n$  dimensions, each vector in this space can be represented by a linear combination of the gradient vectors formed by  $n$  iterations of the conjugate-gradient method. This is true as well for the vector  $\mathbf{m}_0 - \mathbf{m}$ , which points from the solution of equation (1) to the initial model estimate  $\mathbf{m}_0$ . Neglecting computational errors, it takes exactly  $n$  iterations to find this vector by successive optimization of the coefficients. This proves that the conjugate-gradient method converges to the exact solution in a finite number of steps (assuming that the model belongs to a finite-dimensional space).

The method of conjugate gradients simplifies formula (26) to the form

$$\alpha_n = \frac{(\mathbf{r}_{n-1}, \mathbf{A} \mathbf{c}_n)}{\|\mathbf{A} \mathbf{s}_n\|^2} = \frac{(\mathbf{A}^T \mathbf{r}_{n-1}, \mathbf{c}_n)}{\|\mathbf{A} \mathbf{s}_n\|^2} = \frac{\|\mathbf{c}_n\|^2}{\|\mathbf{A} \mathbf{s}_n\|^2}, \quad (31)$$

which in turn leads to the simplification of formula (8), as follows:

$$\|\mathbf{r}_n\|^2 = \|\mathbf{r}_{n-1}\|^2 - \frac{\|\mathbf{c}_n\|^4}{\|\mathbf{A} \mathbf{s}_n\|^2}. \quad (32)$$

If the gradient is not equal to zero, the residual is guaranteed to decrease. If the gradient is equal to zero, we have already found the solution.

## Short memory of the gradients

Substituting the gradient direction (28) into formula (23) and applying formulas (4) and (27), we can see that

$$\beta_n^{(j)} = \frac{(\mathbf{A} \mathbf{c}_n, \mathbf{r}_j - \mathbf{r}_{j-1})}{\alpha_j \|\mathbf{A} \mathbf{s}_j\|^2} = \frac{(\mathbf{c}_n, \mathbf{A}^T \mathbf{r}_j - \mathbf{A}^T \mathbf{r}_{j-1})}{\alpha_j \|\mathbf{A} \mathbf{s}_j\|^2} = \frac{(\mathbf{c}_n, \mathbf{c}_{j+1} - \mathbf{c}_j)}{\alpha_j \|\mathbf{A} \mathbf{s}_j\|^2}. \quad (33)$$

The orthogonality condition (30) and the definition of the coefficient  $\alpha_j$  from equation (31) further transform this formula to the form

$$\beta_n^{(n-1)} = \frac{\|\mathbf{c}_n\|^2}{\alpha_{n-1}\|\mathbf{A}\mathbf{s}_{n-1}\|^2} = \frac{\|\mathbf{c}_n\|^2}{\|\mathbf{c}_{n-1}\|^2}, \quad (34)$$

$$\beta_n^{(j)} = 0, \quad 1 \leq j \leq n-2. \quad (35)$$

Equation (35) shows that the conjugate-gradient method needs to remember only the previous step direction in order to optimize the search at each iteration. This is another remarkable property distinguishing that method in the family of conjugate-direction methods.

## PROGRAM

The program in Table 1 implements one iteration of the conjugate-direction method. It is based upon Jon Claerbout's `cgstep()` program (?) and uses an analogous naming convention. Vectors in the data space are denoted by double letters.

In addition to the previous steps  $\mathbf{s}_j$  and their conjugate counterparts  $\mathbf{A}\mathbf{s}_j$  (array `s`), the program stores the squared norms  $\|\mathbf{A}\mathbf{s}_j\|^2$  (variable `beta`) to avoid recomputation. For practical reasons, the number of remembered iterations can actually be smaller than the total number of iterations.

## EXAMPLES

### Example 1: Inverse interpolation

Matthias Schwab has suggested (in a personal communication) an interesting example, in which the `cgstep` program fails to comply with the conjugate-gradient theory. The inverse problem is a simple one-dimensional data interpolation with a known filter (?). The known portion of the data is a single spike in the middle. One hundred other data points are considered missing. The known filter is the Laplacian  $(1, -2, 1)$ , and the expected result is a bell-shaped cubic spline. The forward problem is strictly linear, and the exact adjoint is easily computed by reverse convolution. However, the conjugate-gradient program requires significantly more than the theoretically predicted 100 iterations. Figure 2 displays the convergence to the final solution in three different plots. According to the figure, the actual number of iterations required for convergence is about 300. Figure 3 shows the result of a similar experiment with the conjugate-direction solver `cdstep`. The number of required iterations is reduced to almost the theoretical one hundred. This indicates that the orthogonality of directions implied in the conjugate-gradient method has been distorted by computational errors. The additional cost of correcting these errors with the conjugate-direction solver comes from storing the preceding 100 directions in memory. A smaller number of memorized steps produces smaller improvements.

```

44 void sf_cdstep(bool forget      /* restart flag */,
45               int nx          /* model size */,
46               int ny          /* data size */,
47               float* x        /* current model [nx] */,
48               const float* g  /* gradient [nx] */,
49               float* rr       /* data residual [ny] */,
50               const float* gg /* conjugate gradient [ny] */)
51 /*< Step of conjugate-direction iteration.
52    The data residual is  $rr = A x - dat$ 
53 >*/
54 {
55     float *s, *si, *ss;
56     double alpha, beta;
57     int i, n, ix, iy;
58
59     s = sf_floatalloc(nx+ny);
60     ss = s+nx;
61
62     for (ix=0; ix < nx; ix++) { s[ix] = g[ix]; }
63     for (iy=0; iy < ny; iy++) { ss[iy] = gg[iy]; }
64
65     sf_llist_rewind (steps);
66     n = sf_llist_depth (steps);
67
68     for (i=0; i < n; i++) {
69         sf_llist_down (steps, &si, &beta);
70         alpha = - cblas_dsdot(ny, gg, 1, si+nx, 1) / beta;
71         cblas_saxpy(nx+ny, alpha, si, 1, s, 1);
72     }
73
74     beta = cblas_dsdot(ny, s+nx, 1, s+nx, 1);
75     if (beta < DBLEPSILON) return;
76
77     sf_llist_add (steps, s, beta);
78     if (forget) sf_llist_chop (steps);
79     alpha = - cblas_dsdot(ny, rr, 1, ss, 1) / beta;
80
81     cblas_saxpy(nx, alpha, s, 1, x, 1);
82     cblas_saxpy(ny, alpha, ss, 1, rr, 1);
83 }

```

Table 1: The source of this program is RSF/api/c/cdstep.c




Figure 2: Convergence of the missing data interpolation problem with the conjugate-gradient solver. Current models are plotted against the number of iterations. The three plots are different displays of the same data.




Figure 3: Convergence of the missing data interpolation problem with the long-memory conjugate-direction solver. Current models are plotted against the number of iterations. The three plots are different displays of the same data.

## Example 2: Velocity transform

The next test example is the velocity transform inversion with a CMP gather from the Mobil AVO dataset (Nichols, 1994; Lumley et al., 1994; Lumley, 1994). I use Jon Claerbout's `veltran` program (Claerbout, 1995b) for anti-aliased velocity transform with rho-filter preconditioning and compare three different pairs of operators for inversion. The first pair is the CMP stacking operator with the "migration" weighting function ( $w = \frac{(t_0/t)}{\sqrt{t}}$ ) and its adjoint. The second pair is the "pseudo-unitary" velocity transform with the weighting proportional to  $\sqrt{|sx|}$ , where  $x$  is the offset and  $s$  is the slowness. These two pairs were used in the velocity transform inversion with the iterative conjugate-gradient solver. The third pair uses the weight proportional to  $|x|$  for CMP stacking and  $|s|$  for the reverse operator. Since these two operators are not exact adjoints, it is appropriate to apply the method of conjugate directions for inversion. The convergence of the three different inversions is compared in Figure 4. We can see that the third method reduces the least-square residual error, though it has a smaller effect than that of the pseudo-unitary weighting in comparison with the uniform one. The results of inversion after 10 conjugate-gradient iterations are plotted in Figures 5 and 6, which are to be compared with the analogous results of Lumley (1994) and Nichols (1994).




Figure 4: Comparison of convergence of the iterative velocity transform inversion. The left plot compares conjugate-gradient inversion with unweighted (uniformly weighted) and pseudo-unitary operators. The right plot compares pseudo-unitary conjugate-gradient and weighted conjugate-direction inversion.




Figure 5: Input CMP gather (left) and its velocity transform counterpart (right) after 10 iterations of conjugate-direction inversion.




Figure 6: The modeled CMP gather (left) and the residual data (right) plotted at the same scale.

## CONCLUSIONS

The conjugate-gradient solver is a powerful method of least-square inversion because of its remarkable algebraic properties. In practice, the theoretical basis of conjugate gradients can be distorted by computational errors. In some applications of inversion, we may want to do that on purpose, by applying inexact adjoints in preconditioning. In both cases, a safer alternative is the method of conjugate directions. Jon Claerbout's `cgstep()` program actually implements a short-memory version of the conjugate-direction method. Extending the length of the memory raises the cost of iterations, but can speed up the convergence.

## REFERENCES

- Claerbout, J., 1995a, Ellipsoids versus hyperboloids, *in* SEP-89: Stanford Exploration Project, 201–206.
- , 2003, Image estimation by example: Geophysical soundings image construction: Multidimensional autoregression: Stanford Exploration Project.
- Claerbout, J. F., 1992, Earth Soundings Analysis: Processing Versus Inversion: Blackwell Scientific Publications.
- , 1995b, Basic Earth Imaging: Stanford Exploration Project.
- Crawley, S., 1995, Approximate vs. exact adjoints in inversion, *in* SEP-89: Stanford Exploration Project, 207–216.
- Fletcher, R., and C. M. Reeves, 1964, Function minimization by conjugate gradients: *Computer Journal*, **7**, 149–154.
- Gill, P. E., W. Murray, and M. H. Wright, 1995, Practical optimization: Academic Press.
- Hestenes, M. R., and E. Stiefel, 1952, Methods of conjugate gradients for solving linear systems: *J. Res. NBS*, **49**, 409–436.
- Kleinman, R. E., and P. M. van den Berg, 1991, Iterative methods for solving integral equations: *Radio Science*, **26**, 175–181.
- Lumley, D., D. Nichols, and T. Rekdal, 1994, Amplitude-preserved multiple suppression, *in* SEP-82: Stanford Exploration Project, 25–45.
- Lumley, D. E., 1994, Estimating a pseudounitary operator for velocity-stack inversion, *in* SEP-82: Stanford Exploration Project, 63–78.
- Nichols, D., 1994, Velocity-stack inversion using  $L_p$  norms, *in* SEP-82: Stanford Exploration Project, 1–16.
- Sevink, A. G. J., and G. C. Herman, 1994, Fast iterative solution of sparsely sampled seismic inverse problems: *Inverse Problems*, **10**, 937–948.