

ALGORITHM 583

LSQR: Sparse Linear Equations and Least Squares Problems

CHRISTOPHER C. PAIGE

McGill University, Canada

and

MICHAEL A. SAUNDERS

Stanford University

Categories and Subject Descriptors: G.1.3 [Numerical Analysis] Numerical Linear Algebra—*linear systems (direct and iterative methods)*, G.3 [Mathematics of Computing]: Probability and Statistics—*statistical computing, statistical software*, G.m [Mathematics of Computing]: Miscellaneous—*FORTRAN program units*

General Terms Algorithms

Additional Key Words and Phrases: Analysis of variance, conjugate-gradient method, least squares, linear equations, regression, sparse matrix

1. INTRODUCTION

LSQR finds a solution x to the following problems:

$$\text{Unsymmetric equations: solve } Ax = b \quad (1.1)$$

$$\text{Linear least squares: minimize } \|Ax - b\|_2 \quad (1.2)$$

$$\text{Damped least squares: minimize } \left\| \begin{bmatrix} A \\ \lambda I \end{bmatrix} x - \begin{bmatrix} b \\ 0 \end{bmatrix} \right\|_2 \quad (1.3)$$

where A is a matrix with m rows and n columns, b is an m -vector, λ is a scalar, and the given data A , b , λ are real. The matrix A will normally be large and sparse. It is defined by means of a user-written subroutine APROD, whose

Received 4 June 1980; revised 23 September 1981, accepted 28 February 1982

This work was supported by Natural Sciences and Engineering Research Council of Canada Grant A8652, by the New Zealand Department of Scientific and Industrial Research; and by U.S. National Science Foundation Grants MCS-7926009 and ECS-8012974, the Department of Energy under Contract AM03-76SF00326, PA No. DE-AT03-76ER72018, the Office of Naval Research under Contract N00014-75-C-0267, and the Army Research Office under Contract DAA29-79-C-0110.

Authors' addresses: C. C. Paige, School of Computer Science, McGill University, Montreal, Quebec, Canada H3A 2K6; M. A. Saunders, Department of Operations Research, Stanford University, Stanford, CA 94305.

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association for Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission.

© 1982 ACM 0098-3500/82/0600-0195 \$00 75

Table I. Comparison of CGLS and LSQR

	Storage	Work per iteration
CGLS, $\lambda = 0$	$2m + 2n$	$2m + 3n$
CGLS, $\lambda \neq 0$	$2m + 2n$	$2m + 5n$
LSQR, any λ	$m + 2n$	$3m + 5n$

essential function is to compute products of the form Ax and $A^T y$ for given vectors x and y .

Problems (1.1) and (1.2) are treated as special cases of (1.3), which we shall write as

$$\min \|\bar{A}x - \bar{b}\|_2, \quad \bar{A} = \begin{bmatrix} A \\ \lambda I \end{bmatrix}, \quad \bar{b} = \begin{bmatrix} b \\ 0 \end{bmatrix}. \quad (1.4)$$

An earlier successful method for such problems is the *conjugate-gradient method for least squares systems* given by Hestenes and Stiefel [3]. (This method is described as algorithm CGLS in [6, sect. 7.1].) CGLS and LSQR are iterative methods with similar qualitative properties. Their computational requirements are summarized in Table I. In addition they require a product Ax and a product $A^T y$ each iteration.

In order to achieve the storage shown for LSQR, we ask the user to implement the matrix-vector products in the form

$$y \leftarrow y + Ax \quad \text{and} \quad x \leftarrow x + A^T y, \quad (1.5)$$

where \leftarrow means that one of the given vectors is overwritten by the expression shown. (A parameter specifies which expression the user's subroutine APROD should compute on any given entry.) We see that LSQR has a storage advantage if the operations (1.5) can be performed with no additional storage beyond that required to represent A . For least squares applications with many observations ($m \gg n$), this could be useful.

The work shown in Table I is the number of floating-point multiplications per iteration, excluding the work involved in the products Ax , $A^T y$. Since CGLS is somewhat more efficient, we would not discourage using that method whenever A or \bar{A} is well conditioned. However, LSQR is likely to obtain a more accurate solution in fewer iterations if \bar{A} is moderately or severely ill-conditioned.

Let $\bar{r}_k = \bar{b} - \bar{A}x_k$ be the residual vector associated with the k th iteration. LSQR provides estimates of $\|x_k\|_2$, $\|\bar{r}_k\|_2$, $\|\bar{A}^T \bar{r}_k\|_2$, the norm of \bar{A} , the condition number of \bar{A} , and standard errors for the components of x . The last two items require a further $2n$ multiplications per iteration and an additional n -vector of storage.

Subroutine LSQR is written in the PFORT subset of American National Standard FORTRAN. It contains no machine-dependent constants. Auxiliary routines required are APROD, NORMLZ, SCOPY, SNRM2, and SSCAL. The last three correspond to members of the BLAS collection [5].

2. MATHEMATICAL BACKGROUND

Algorithmic details are given in [6], mainly for the case $\lambda = 0$. We summarize these here with λ reintroduced, and show that a given value of λ may be dealt with at negligible cost. The vector norm $\|v\|_2 = (v^T v)^{1/2}$ is used throughout.

LSQR uses an algorithm of Golub and Kahan to reduce A to lower bidiagonal form. The quantities produced from A and b after $k + 1$ steps of the bidiagonalization (procedure Bidiag 1 [6]) are

$$\begin{aligned} \beta_1 &= \|b\|, \\ U_{k+1} &= [u_1, u_2, \dots, u_{k+1}], \\ V_{k+1} &= [v_1, v_2, \dots, v_{k+1}], \end{aligned} \quad B_k = \begin{bmatrix} \alpha_1 & & & & & \\ \beta_2 & \alpha_2 & & & & \\ & \beta_3 & \alpha_3 & & & \\ & & \ddots & \ddots & & \\ & & & \ddots & \alpha_k & \\ & & & & & \beta_{k+1} \end{bmatrix}. \quad (2.1)$$

The k th approximation to the solution x is then defined to be $x_k = V_k y_k$, where y_k solves the subproblem

$$\min \left\| \begin{bmatrix} B_k \\ \lambda I \end{bmatrix} y_k - \begin{bmatrix} \beta_1 e_1 \\ 0 \end{bmatrix} \right\|. \quad (2.2)$$

Letting the associated residual vectors be

$$\begin{aligned} t_{k+1} &= \beta_1 e_1 - B_k y_k \\ r_k &= b - A x_k \\ \bar{r}_k &= \bar{b} - \bar{A} x_k, \end{aligned} \quad (2.3)$$

we find that the relations

$$\begin{aligned} r_k &= U_{k+1} t_{k+1} \\ A^T r_k &= \lambda^2 x_k + \alpha_{k+1} \tau_{k+1} U_{k+1} \end{aligned} \quad (2.4)$$

will hold to machine accuracy, where τ_{k+1} is the last component of t_{k+1} , and we therefore conclude that (r_k, x_k) will be an acceptable solution of (1.4) if the computed value of either $\|t_{k+1}\|$ or $|\alpha_{k+1} \tau_{k+1}|$ is suitably small.

Bjorck [1] has previously observed that subproblem (2.2) is the appropriate generalization of $\min \|B_k y_k - \beta_1 e_1\|$, when $\lambda \neq 0$. He also discusses methods for computing y_k and x_k efficiently for various λ and k .

In LSQR we assume that a single value of λ is given, and to save storage and work, we do not compute $y_k, r_k,$ or t_{k+1} . The orthogonal factorization

$$Q_k \begin{bmatrix} B_k & \beta_1 e_1 \\ \lambda I & 0 \end{bmatrix} = \begin{bmatrix} R_k & f_k \\ 0 & \bar{\Phi}_{k+1} \\ 0 & q_k \end{bmatrix} \quad (2.5)$$

is computed ($Q_k^T Q_k = I; R_k$ upper bidiagonal, $k \times k$) and this would give $R_k y_k = f_k$, but instead we solve $R_k^T D_k^T = V_k^T$ and form $x_k = D_k f_k$.

The factorization (2.5) is formed similarly to the case $\lambda = 0$ in [6], except that two rotations are required per step instead of one. For $k = 2$, the factorization

proceeds according to

$$\begin{aligned}
 & \begin{bmatrix} \alpha_1 & & & \beta_1 \\ \beta_2 & \alpha_2 & & \\ & \beta_3 & & \\ \lambda & & & \lambda \end{bmatrix} \rightarrow \begin{bmatrix} \tilde{\rho}_1 & & & \tilde{\phi}_1 \\ \beta_2 & \alpha_2 & & \\ & \beta_3 & & \\ & & & \psi_1 \\ & & & \lambda \end{bmatrix} \rightarrow \begin{bmatrix} \rho_1 & \theta_2 & \phi_1 \\ & \tilde{\rho}_2 & \tilde{\phi}_2 \\ & \beta_3 & \\ & & \psi_1 \\ & & & \lambda \end{bmatrix} \\
 & \rightarrow \begin{bmatrix} \rho_1 & \theta_2 & \phi_1 \\ & \tilde{\rho}_2 & \tilde{\phi}_2 \\ & \beta_3 & \\ & & \psi_1 \\ & & & \psi_2 \end{bmatrix} \rightarrow \begin{bmatrix} \rho_1 & \theta_2 & \phi_1 \\ & \rho_2 & \phi_2 \\ & & \tilde{\phi}_3 \\ & & \psi_1 \\ & & & \psi_2 \end{bmatrix} .
 \end{aligned}$$

Note that the first λ is rotated into the diagonal element α_1 . This alters the right-hand side $\beta_1 e_1$ to produce ψ_1 , the first component of q_k . An alternative is to rotate λ into β_2 (and similarly for later λ), since this does not affect the right-hand side and it more closely simulates the algorithm that results when LSQR is applied to \bar{A} and \bar{b} directly. However, the rotations then have a greater effect on B_k , and in practice the first option has proved to give marginally more accurate results.

The estimates required to implement the stopping criteria are

$$\begin{aligned}
 \|\bar{r}_k\|^2 &= \|r_k\|^2 + \lambda^2 \|x_k\|^2 \approx \bar{\phi}_{k+1}^2 + \|q_k\|^2, \\
 \|\bar{A}^T \bar{r}_k\| &= \|A^T r_k - \lambda^2 x_k\| \approx \left| \frac{\alpha_{k+1} \beta_{k+1} \phi_k}{\rho_k} \right|.
 \end{aligned}$$

This is a simple generalization of the case $\lambda = 0$. No additional storage is needed for q_k , since only its norm is required. In short, although the presence of λ complicates the algorithm description, it adds essentially nothing to the storage and work per iteration.

3. REGULARIZATION AND RELATED WORK

Introducing λ as in (1.3) is just one way of “regularizing” the solution x , in the sense that it can reduce the size of the computed solution and make its components less sensitive to changes in the data. LSQR is applicable when a value of λ is known a priori. The value is entered via the subroutine parameter DAMP. A second method for regularizing x is available through LSQR’s parameter ACOND, which can cause iterations to terminate before $\|x_k\|$ becomes large. A similar approach has recently been described by Wold et al. [9], who give an illuminating interpretation of the bidiagonalization as a partial least squares procedure. Their description will also be useful to those who prefer the notation of multiple regression.

Methods for choosing λ , and other approaches to regularization, are given in [1, 2, 4, 8] and elsewhere. For a philosophical discussion, see [7].

4. CODING APROD

The best way to compute $y + Ax$ and $x + A^T y$ depends upon the origin of the matrix A . We shall illustrate a case that commonly arises, in which A is a sparse matrix whose nonzero coefficients are stored by rows in a simple list. Let A have

M rows, N columns, and NZ nonzeros. Conceptually we need three arrays dimensioned as REAL $RA(NZ)$ and INTEGER $JA(NZ)$, $NA(M)$, where

$RA(L)$ is the L th nonzero of A , counting across row 1, then across row 2, and so on;

$JA(L)$ is the column in which the L th nonzero of A lies;

$NA(I)$ is the number of nonzero coefficients in the I th row of A .

These quantities may be used in a straightforward way, as shown in Figure 1 (a FORTRAN implementation). We assume that they are made available to APROD through COMMON, and that the actual array dimensions are suitably large.

Blank or labeled COMMON will often be convenient for transmitting data to APROD. (Of course, some of the data could be local to APROD.) For greater generality, the parameter lists for LSQR and APROD include two workspace arrays IW, RW and their lengths LENIW, LENRW. LSQR does not use these parameters directly; it just passes them to APROD.

Figure 2 illustrates their use on the same example (sparse A stored by rows). An auxiliary subroutine APROD1 is needed to make the code readable. A similar scheme should be used to initialize the workspace parameters prior to calling LSQR.

Returning to the example itself, it may often be natural to store A by *columns* rather than rows, using analogous data structures. However, we note that in sparse least squares applications, A may have many more rows than columns ($M \gg N$). In such cases it is vital to store A by rows as shown, if the machine being used has a paged (virtual) memory. Random access is then restricted to arrays of length N rather than M , and page faults will therefore be kept to a minimum.

Note also that the arrays RA , JA , NA are adequate for computing both Ax and $A^T y$; we do not need to store A by rows *and* by columns.

Regardless of the application, it will be apparent when coding APROD for the two values of MODE that the matrix A is effectively being defined *twice*. Great care must be taken to avoid coding inconsistent expressions $y + A_1 x$ and $x + A_2^T y$, where either A_1 or A_2 is different from the desired A . (If $A_1 \neq A_2$, algorithm LSQR will not converge.) Parameters ANORM, ACOND, and CONLIM provide a safeguard for such an event.

5. PRECONDITIONING

It is well known that conjugate-gradient methods can be accelerated if a nonsingular matrix M is available to approximate A in some useful sense. When A is square and nonsingular, the system $Ax = b$ is equivalent to both of the following systems:

$$(M^{-1}A)x = c \quad \text{where } Mc = b; \quad (5.1)$$

$$(AM^{-1})z = b \quad \text{where } Mx = z. \quad (5.2)$$

For least squares systems (undamped), only the analogue of (5.2) is applicable:

$$\min \|Ax - b\|_2 = \min \|(AM^{-1})z - b\|_2, \quad \text{where } Mx = z. \quad (5.3)$$

```

SUBROUTINE APROD( MODE,M,N,X,Y,
*                LENIW,LENRW,IW,RW )
C
C   INTEGER    MODE,M,N,LENIW,LENRW
C   INTEGER    IW(LENIW)
C   REAL       X(N),Y(M),RW(LENRW)
C
C   APROD PERFORMS THE FOLLOWING FUNCTIONS:
C
C       IF MODE = 1, SET Y = Y + A*X
C       IF MODE = 2, SET X = X + A(TRANSPPOSE)*Y
C
C   WHERE A IS A MATRIX STORED BY ROWS IN
C   THE ARRAYS RA, JA, NA. IN THIS EXAMPLE,
C   RA, JA, NA ARE STORED IN COMMON.
C
C   REAL       RA
C   INTEGER    JA,NA
C   COMMON     RA(9000),JA(9000),NA(1000)
C
C   INTEGER    I,J,L,L1,L2
C   REAL       SUM,YI,ZERO
C
C   ZERO = 0.0
C   L2 = 0
C   IF (MODE .NE.1) GO TO 400
C
C   -----
C   MODE = 1 -- SET Y = Y + A*X.
C   -----
C   DO 200 I = 1, M
C       SUM = ZERO
C       L1 = L2 + 1
C       L2 = L2 + NA(I)
C       DO 100 L = L1, L2
C           J = JA(L)
C           SUM = SUM + RA(L)*X(J)
C   100  CONTINUE
C       Y(I) = Y(I) + SUM
C   200  CONTINUE
C       RETURN
C
C   -----
C   MODE = 2 -- SET X = X + A(TRANSPPOSE)*Y.
C   -----
C   400 DO 600 I = 1, M
C       YI = Y(I)
C       L1 = L2 + 1
C       L2 = L2 + NA(I)
C       DO 500 L = L1, L2
C           J = JA(L)
C           X(J) = X(J) + RA(L)*YI
C   500  CONTINUE
C   600  CONTINUE
C       RETURN
C
C   END OF APROD
C   END

```

Fig. 1. Computation of $y + Ax$, $x + A^T y$, where A is a sparse matrix stored compactly by rows. For convenience, the data structure for A is held in COMMON.

```

SUBROUTINE APROD( MODE,M,N,X,Y,
*           LENIW,LENRW,IW,RW )
C
C   INTEGER    MODE,M,N,LENIW,LENRW
C   INTEGER    IW(LENIW)
C   REAL       X(N),Y(M),RW(LENRW)
C
C   APROD PERFORMS THE FOLLOWING FUNCTIONS:
C
C       IF MODE = 1, SET  Y = Y + A*X
C       IF MODE = 2, SET  X = X + A(TRANSPPOSE)*Y
C
C   WHERE A IS A MATRIX STORED BY ROWS IN
C   THE ARRAYS RA, JA, NA. IN THIS EXAMPLE,
C   APROD IS AN INTERFACE BETWEEN LSQR AND
C   ANOTHER USER ROUTINE THAT DOES THE WORK.
C   THE WORKSPACE ARRAY RW CONTAINS RA.
C   THE FIRST M COMPONENTS OF IW CONTAIN NA,
C   AND THE REMAINDER OF IW CONTAINS JA.
C   THE DIMENSIONS OF RW AND IW ARE ASSUMED
C   TO BE SUFFICIENTLY LARGE.
C
C   INTEGER    LENJA,LENRA,LOCJA
C
C   LOCJA = M + 1
C   LENJA = LENIW - LOCJA + 1
C   LENRA = LENRW
C   CALL APROD1( MODE,M,N,X,Y,
*           LENJA,LENRA,IW,IW(LOCJA),RW )
C   RETURN
C
C   END OF APROD
C   END

SUBROUTINE APROD1( MODE,M,N,X,Y,
*           LENJA,LENRA,NA,JA,RA )
C
C   INTEGER    MODE,M,N,LENJA,LENRA
C   INTEGER    NA(M),JA(LENJA)
C   REAL       X(N),Y(M),RA(LENRA)
C
C   APROD1 DOES THE WORK FOR APROD.
C
C   INTEGER    I,J,L,L1,L2
C   REAL       SUM,YI,ZERO
C
C   < the same code as in APROD in Figure 1 >
C
C   END OF APROD1
C   END

```

Fig. 2 Same as Figure 1, with the data structure for A held in the workspace parameters.

LEAST-SQUARES TEST PROBLEM P(20 10 1 1 1.00E-03)
 CONDITION NO. = 9.9995E 00 RESIDUAL FUNCTION = 9.812157000E-01

LSQR -- LEAST-SQUARES SOLUTION OF A*X = B
 THE MATRIX A HAS 20 ROWS AND 10 COLS
 THE DAMPING PARAMETER IS DAMP = 1.00E-03
 ATOL = 1.00E-06 CONLM = 1.00E 02
 BTOL = 1.00E-06 ITNLM = 80

ITN	X(1)	FUNCTION	COMPATIBLE INCOMPATIBLE NORM(ABAR)	COND(ABAR)
0	0.0000000000E-01	6.3410580000E 00	1.000E 00	9.135E-02
1	-6.3564250000E-01	4.0387670000E 00	6.369E-01	7.244E-01
2	-4.7282630000E-01	2.3303970000E 00	3.675E-01	3.349E-01
3	-2.8075080000E-01	1.8962160000E 00	2.990E-01	2.462E-01
4	2.6825070000E-01	1.5905200000E 00	2.508E-01	1.424E-01
5	1.2649560000E 00	1.4032960000E 00	2.213E-01	1.160E-01
6	2.0648040000E 00	1.2910880000E 00	2.036E-01	9.273E-02
7	3.0031450000E 00	1.2072930000E 00	1.904E-01	8.294E-02
8	3.7526340000E 00	1.1551220000E 00	1.822E-01	5.138E-02
9	5.4443550000E 00	1.0780640000E 00	1.700E-01	3.155E-02
10	8.9918140000E 00	9.8151740000E-01	1.548E-01	1.283E-02
11	8.9998990000E 00	9.8120520000E-01	1.547E-01	2.017E-04
12	8.9999290000E 00	9.8120540000E-01	1.547E-01	4.361E-06
13	8.9999280000E 00	9.8120550000E-01	1.547E-01	9.078E-07
				2.20E 00
				2.38E 00
				2.48E 00
				2.75E 01
				2.98E 01
				3.13E 01

NO. OF ITERATIONS = 13 STOPPING CONDITION = 2

THE LEAST-SQRS SOLN IS GOOD ENOUGH, GIVEN ATOL

	RESIDUAL NORM (ABAR*X - BBAR)	RESIDUAL NORM (NORMAL EQNS)	SOLUTION NORM (X)
ESTIMATED BY LSQR	9.812055E-01	2.206693E-06	1.688187E 01
COMPUTED FROM X	9.812157E-01	1.083419E-05	1.688184E 01

SOLUTION			
1 8.99993	2 7.99997	3 6.99998	4 5.99999
6 3.99999	7 3.00000	8 2.00000	9 0.999994
			10 -0.204206E-05
STANDARD ERRORS			
1 2.11589	2 0.888101	3 0.685644	4 0.556184
6 0.409182	7 0.565480	8 0.519385	9 0.375466
			10 0.589787

Fig. 3. Example output from test program and LSQR on a damped least squares problem.

We note only that *subroutine LSQR may be applied without change* to systems (5.1)–(5.3). The effect of M is localized to the user's own subroutine APROD. For example, when $\text{MODE} = 1$, APROD for the last two systems should compute $y + (AM^{-1})x$ by first solving $Mw = x$ and then computing $y + Aw$. Clearly it must be possible to solve systems involving M and M^T very efficiently.

6. OUTPUT

Subroutine LSQR produces printed output on file NOUT, if the parameter NOUT is positive. This is illustrated in Figure 3, in which the least squares problem solved is $P(20, 10, 1, 1)$ as defined in [6], with a slight generalization to include a damping parameter $\lambda = 10^{-3}$. (Single precision was used on an IBM 370/168.) The items printed at the k th iteration are as follows.

ITN	The iteration number k . Results are always printed for the first 10 and last 10 iterations. Intermediate results are printed if $m \leq 40$ or $n \leq 40$, or if one of the convergence conditions is nearly satisfied. Otherwise, information is printed every 10th iteration.
X(1)	The value of the first element of the approximate solution x_k .
FUNCTION	The value of the function being minimized, namely $\ \bar{r}_k\ = (\ r_k\ ^2 + \lambda^2 \ x_k\ ^2)^{1/2}$.
COMPATIBLE	A dimensionless quantity which should converge to zero <i>if and only if</i> $Ax = b$ is compatible. It is an estimate of $\ \bar{r}_k\ / \ b\ $, which decreases monotonically.
INCOMPATIBLE	A dimensionless quantity which should converge to zero <i>if and only if</i> the optimum $\ \bar{r}_k\ $ is nonzero. It is an estimate of $\ \bar{A}^T \bar{r}_k\ / (\ \bar{A}\ _F \ \bar{r}_k\)$, which is usually <i>not</i> monotonic.
NORM(ABAR)	A monotonically increasing estimate of $\ \bar{A}\ _F$.
COND(ABAR)	A monotonically increasing estimate of $\text{cond}(\bar{A}) = \ \bar{A}\ _F \ \bar{A}^+\ _F$, the condition number of \bar{A} .

ACKNOWLEDGMENT

The authors are grateful to Richard Hanson for suggestions that prompted several improvements to the implementation of LSQR.

REFERENCES

1. BJORCK, Å. A bidiagonalization algorithm for solving ill-posed systems of linear equations Rep LITH-MAT-R-80-33, Dep. Mathematics, Linköping Univ., Linköping, Sweden, 1980.
2. ELDÉN, L. Algorithms for the regularization of ill-conditioned least squares problems *BIT* 17 (1977), 134–145.
3. HESTENES, M.R., AND STIEFEL, E. Methods of conjugate gradients for solving linear systems. *J Res. N.B.S.* 49 (1952), 409–436.
4. LAWSON, C.L., AND HANSON, R.J. *Solving Least Squares Problems*. Prentice-Hall, Englewood Cliffs, N.J., 1974.
5. LAWSON, C.L., HANSON, R.J., KINCAID, D.R., AND KROGH, F.T. Basic linear algebra subprograms for Fortran usage *ACM Trans Math Softw* 5, 3 (Sept 1979), 308–323 and (Algorithm) 324–325.
6. PAIGE, C.C., AND SAUNDERS, M.A. LSQR An algorithm for sparse linear equations and sparse least squares *ACM Trans Math Softw* 8, 1 (March 1982), 43–71

7. SMITH, G., AND CAMPBELL, F. A critique of some ridge regression methods. *J. Am. Stat. Assoc.* 75, 369 (March 1980), 74-81
8. VARAH, J.M. A practical examination of some numerical methods for linear discrete ill-posed problems. *SIAM Rev.* 21 (1979), 100-111.
9. WOLD, S., WOLD, H., DUNN, W.J., AND RUHE, A. The collinearity problem in linear and nonlinear regression. The partial least squares (PLS) approach to generalized inverses. Rep. UMINF-83.80, Univ. Umeå, Umeå, Sweden, 1980.

ALGORITHM

[A part of the listing is printed here. The complete listing is available from the ACM Algorithms Distribution Service (see page 227 for order form).]

```

SUBROUTINE LSQR( M,N,APROD,DAMP,                                1.
1      LENIW,LENRW,IW,RW,                                     2.
2      U,V,W,X,SE,                                           3.
3      ATOL,BTOL,CONLIM,ITNLIM,NOUT,                         4.
4      ISTOP,ANORM,ACOND,RNORM,ARNORM,XNORM )                 5.
C                                                                 6.
EXTERNAL  APROD                                              7.
INTEGER  M,N,LENIW,LENRW,ITNLIM,NOUT,ISTOP                  8.
INTEGER  IW(LENIW)                                          9.
REAL     RW(LENRW),U(M),V(N),W(N),X(N),SE(N),             10.
1      ATOL,BTOL,CONLIM,DAMP,ANORM,ACOND,RNORM,ARNORM,XNORM 11.
-----                                                    12.
C                                                                 13.
C  LSQR FINDS A SOLUTION X TO THE FOLLOWING PROBLEMS...    14.
C                                                                 15.
C  1. UNSYMMETRIC EQUATIONS -- SOLVE A*X = B                16.
C                                                                 17.
C  2. LINEAR LEAST SQUARES -- SOLVE A*X = B                18.
C     IN THE LEAST-SQUARES SENSE                           19.
C                                                                 20.
C  3. DAMPED LEAST SQUARES -- SOLVE ( A ) * X = ( B )      21.
C     ( DAMP*I ) ( 0 )                                     22.
C     IN THE LEAST-SQUARES SENSE                           23.
C                                                                 24.
C  WHERE A IS A MATRIX WITH M ROWS AND N COLUMNS, B IS AN 25.
C  M-VECTOR, AND DAMP IS A SCALAR (ALL QUANTITIES REAL).    26.
C  THE MATRIX A IS INTENDED TO BE LARGE AND SPARSE. IT IS 27.
C  ACCESSED BY MEANS OF SUBROUTINE CALLS OF THE FORM        28.
C                                                                 29.
C     CALL APROD( MODE,M,N,X,Y,LENIW,LENRW,IW,RW )          30.
C                                                                 31.
C  WHICH MUST PERFORM THE FOLLOWING FUNCTIONS...            32.
C                                                                 33.
C     IF MODE = 1, COMPUTE Y = Y + A*X.                      34.
C     IF MODE = 2, COMPUTE X = X + A(TRANPOSE)*Y.           35.
C                                                                 36.
C  THE VECTORS X AND Y ARE INPUT PARAMETERS IN BOTH CASES.  37.
C  IF MODE = 1, Y SHOULD BE ALTERED WITHOUT CHANGING X.    38.
C  IF MODE = 2, X SHOULD BE ALTERED WITHOUT CHANGING Y.    39.
C  THE PARAMETERS LENIW, LENRW, IW, RW MAY BE USED FOR      40.
C  WORKSPACE AS DESCRIBED BELOW.                            41.
C                                                                 42.
C  THE RHS VECTOR B IS INPUT VIA U, AND SUBSEQUENTLY        43.
C  OVERWRITTEN.                                             44.
C                                                                 45.

```

C NOTE. LSQR USES AN ITERATIVE METHOD TO APPROXIMATE THE SOLUTION. 46.
 C THE NUMBER OF ITERATIONS REQUIRED TO REACH A CERTAIN ACCURACY 47.
 C DEPENDS STRONGLY ON THE SCALING OF THE PROBLEM. POOR SCALING OF 48.
 C THE ROWS OR COLUMNS OF A SHOULD THEREFORE BE AVOIDED WHERE 49.
 C POSSIBLE. 50.
 C 51.
 C FOR EXAMPLE, IN PROBLEM 1 THE SOLUTION IS UNALTERED BY 52.
 C ROW-SCALING. IF A ROW OF A IS VERY SMALL OR LARGE COMPARED TO 53.
 C THE OTHER ROWS OF A, THE CORRESPONDING ROW OF (A B) SHOULD 54.
 C BE SCALED UP OR DOWN. 55.
 C 56.
 C IN PROBLEMS 1 AND 2, THE SOLUTION X IS EASILY RECOVERED 57.
 C FOLLOWING COLUMN-SCALING. IN THE ABSENCE OF BETTER INFORMATION, 58.
 C THE NONZERO COLUMNS OF A SHOULD BE SCALED SO THAT THEY ALL HAVE 59.
 C THE SAME EUCLIDEAN NORM (E.G. 1.0). 60.
 C 61.
 C IN PROBLEM 3, THERE IS NO FREEDOM TO RE-SCALE IF DAMP IS 62.
 C NONZERO. HOWEVER, THE VALUE OF DAMP SHOULD BE ASSIGNED ONLY 63.
 C AFTER ATTENTION HAS BEEN PAID TO THE SCALING OF A. 64.
 C 65.
 C THE PARAMETER DAMP IS INTENDED TO HELP REGULARIZE 66.
 C ILL-CONDITIONED SYSTEMS, BY PREVENTING THE TRUE SOLUTION FROM 67.
 C BEING VERY LARGE. ANOTHER AID TO REGULARIZATION IS PROVIDED BY 68.
 C THE PARAMETER ACOND, WHICH MAY BE USED TO TERMINATE ITERATIONS 69.
 C BEFORE THE COMPUTED SOLUTION BECOMES VERY LARGE. 70.
 C 71.
 C 72.
 C NOTATION 73.
 C ----- 74.
 C 75.
 C THE FOLLOWING QUANTITIES ARE USED IN DISCUSSING THE SUBROUTINE 76.
 C PARAMETERS... 77.
 C 78.
 C ABAR = (A), BBAR = (B) 79.
 C (DAMP*I) (0) 80.
 C 81.
 C R = B - A*X, RBAR = BBAR - ABAR*X 82.
 C 83.
 C RNORM = Sqrt(NORM(R)**2 + DAMP**2 * NORM(X)**2) 84.
 C = NORM(RBAR) 85.
 C 86.
 C RELPR = THE RELATIVE PRECISION OF FLOATING-POINT ARITHMETIC 87.
 C ON THE MACHINE BEING USED. FOR EXAMPLE, ON THE IBM 370, 88.
 C RELPR IS ABOUT 1.0E-6 AND 1.0D-16 IN SINGLE AND DOUBLE 89.
 C PRECISION RESPECTIVELY. 90.
 C 91.
 C LSQR MINIMIZES THE FUNCTION RNORM WITH RESPECT TO X. 92.
 C 93.
 C 94.
 C PARAMETERS 95.
 C ----- 96.
 C 97.
 C M INPUT THE NUMBER OF ROWS IN A. 98.
 C 99.
 C N INPUT THE NUMBER OF COLUMNS IN A. 100.
 C 101.
 C APROD EXTERNAL SEE ABOVE. 102.
 C 103.
 C DAMP INPUT THE DAMPING PARAMETER FOR PROBLEM 3 ABOVE. 104.
 C (DAMP SHOULD BE 0.0 FOR PROBLEMS 1 AND 2.) 105.

C IF THE SYSTEM $A \cdot X = B$ IS INCOMPATIBLE, VALUES 106.
 C OF DAMP IN THE RANGE \emptyset TO $\text{SQRT}(\text{RELPR}) \cdot \text{NORM}(A)$ 107.
 C WILL PROBABLY HAVE A NEGLIGIBLE EFFECT. 108.
 C LARGER VALUES OF DAMP WILL TEND TO DECREASE 109.
 C THE NORM OF X AND TO REDUCE THE NUMBER OF 110.
 C ITERATIONS REQUIRED BY LSQR. 111.
 C 112.
 C THE WORK PER ITERATION AND THE STORAGE NEEDED 113.
 C BY LSQR ARE THE SAME FOR ALL VALUES OF DAMP. 114.
 C 115.
 C LENIW INPUT THE LENGTH OF THE WORKSPACE ARRAY IW. 116.
 C LENRW INPUT THE LENGTH OF THE WORKSPACE ARRAY RW. 117.
 C IW WORKSPACE AN INTEGER ARRAY OF LENGTH LENIW. 118.
 C RW WORKSPACE A REAL ARRAY OF LENGTH LENRW. 119.
 C 120.
 C NOTE. LSQR DOES NOT EXPLICITLY USE THE PREVIOUS FOUR 121.
 C PARAMETERS, BUT PASSES THEM TO SUBROUTINE APROD FOR 122.
 C POSSIBLE USE AS WORKSPACE. IF APROD DOES NOT NEED 123.
 C IW OR RW, THE VALUES $\text{LENIW} = 1$ OR $\text{LENRW} = 1$ SHOULD 124.
 C BE USED, AND THE ACTUAL PARAMETERS CORRESPONDING TO 125.
 C IW OR RW MAY BE ANY CONVENIENT ARRAY OF SUITABLE TYPE. 126.
 C 127.
 C U(M) INPUT THE RHS VECTOR B. BEWARE THAT U IS 128.
 C OVER-WRITTEN BY LSQR. 129.
 C 130.
 C V(N) WORKSPACE 131.
 C W(N) WORKSPACE 132.
 C 133.
 C X(N) OUTPUT RETURNS THE COMPUTED SOLUTION X. 134.
 C 135.
 C SE(N) OUTPUT RETURNS STANDARD ERROR ESTIMATES FOR THE 136.
 C COMPONENTS OF X. FOR EACH I, SE(I) IS SET 137.
 C TO THE VALUE $\text{RNORM} \cdot \text{SQRT}(\text{SIGMA}(I,I) / T)$, 138.
 C WHERE $\text{SIGMA}(I,I)$ IS AN ESTIMATE OF THE I-TH 139.
 C DIAGONAL OF THE INVERSE OF $\text{ABAR}(\text{TRANPOSE}) \cdot \text{ABAR}$ 140.
 C AND $T = 1$ IF $M \leq N$, 141.
 C $T = M - N$ IF $M > N$ AND $\text{DAMP} = \emptyset$, 142.
 C $T = M$ IF $\text{DAMP} \neq \emptyset$. 143.
 C 144.
 C ATOL INPUT AN ESTIMATE OF THE RELATIVE ERROR IN THE DATA 145.
 C DEFINING THE MATRIX A. FOR EXAMPLE, 146.
 C IF A IS ACCURATE TO ABOUT 6 DIGITS, SET 147.
 C $\text{ATOL} = 1.0\text{E}-6$. 148.
 C 149.
 C BTOL INPUT AN ESTIMATE OF THE RELATIVE ERROR IN THE DATA 150.
 C DEFINING THE RHS VECTOR B. FOR EXAMPLE, 151.
 C IF B IS ACCURATE TO ABOUT 6 DIGITS, SET 152.
 C $\text{BTOL} = 1.0\text{E}-6$. 153.
 C 154.
 C CONLIM INPUT AN UPPER LIMIT ON $\text{COND}(\text{ABAR})$, THE APPARENT 155.
 C CONDITION NUMBER OF THE MATRIX ABAR. 156.
 C ITERATIONS WILL BE TERMINATED IF A COMPUTED 157.
 C ESTIMATE OF $\text{COND}(\text{ABAR})$ EXCEEDS CONLIM. 158.
 C THIS IS INTENDED TO PREVENT CERTAIN SMALL OR 159.
 C ZERO SINGULAR VALUES OF A OR ABAR FROM 160.
 C COMING INTO EFFECT AND CAUSING UNWANTED GROWTH 161.
 C IN THE COMPUTED SOLUTION. 162.
 C 163.
 C CONLIM AND DAMP MAY BE USED SEPARATELY OR 164.
 C TOGETHER TO REGULARIZE ILL-CONDITIONED SYSTEMS. 165.

C			166.
C		NORMALLY, CONLIM SHOULD BE IN THE RANGE	167.
C		1000 TO 1/RELPR.	168.
C		SUGGESTED VALUE --	169.
C		CONLIM = 1/(100*RELPR) FOR COMPATIBLE SYSTEMS,	170.
C		CONLIM = 1/(10*SQRT(RELPR)) FOR LEAST SQUARES.	171.
C			172.
C		NOTE. IF THE USER IS NOT CONCERNED ABOUT THE PARAMETERS	173.
C		ATOL, BTOL AND CONLIM, ANY OR ALL OF THEM MAY BE SET	174.
C		TO ZERO. THE EFFECT WILL BE THE SAME AS THE VALUES	175.
C		RELPR, RELPR AND 1/RELPR RESPECTIVELY.	176.
C			177.
C	ITNLIM	INPUT AN UPPER LIMIT ON THE NUMBER OF ITERATIONS.	178.
C		SUGGESTED VALUE --	179.
C		ITNLIM = N/2 FOR WELL CONDITIONED SYSTEMS,	180.
C		ITNLIM = 4*N OTHERWISE.	181.
C			182.
C	NOUT	INPUT FILE NUMBER FOR PRINTER. IF POSITIVE,	183.
C		A SUMMARY WILL BE PRINTED ON FILE NOUT.	184.
C			185.
C	ISTOP	OUTPUT AN INTEGER GIVING THE REASON FOR TERMINATION...	186.
C			187.
C		0 X = 0 IS THE EXACT SOLUTION.	188.
C		NO ITERATIONS WERE PERFORMED.	189.
C			190.
C		1 THE EQUATIONS A*X = B ARE PROBABLY	191.
C		COMPATIBLE. NORM(A*X - B) IS SUFFICIENTLY	192.
C		SMALL, GIVEN THE VALUES OF ATOL AND BTOL.	193.
C			194.
C		2 THE SYSTEM A*X = B IS PROBABLY NOT	195.
C		COMPATIBLE. A LEAST-SQUARES SOLUTION HAS	196.
C		BEEN OBTAINED WHICH IS SUFFICIENTLY ACCURATE,	197.
C		GIVEN THE VALUE OF ATOL.	198.
C			199.
C		3 AN ESTIMATE OF COND(ABAR) HAS EXCEEDED	200.
C		CONLIM. THE SYSTEM A*X = B APPEARS TO BE	201.
C		ILL-CONDITIONED. OTHERWISE, THERE COULD BE AN	202.
C		AN ERROR IN SUBROUTINE APROD.	203.
C			204.
C		4 THE EQUATIONS A*X = B ARE PROBABLY	205.
C		COMPATIBLE. NORM(A*X - B) IS AS SMALL AS	206.
C		SEEMS REASONABLE ON THIS MACHINE.	207.
C			208.
C		5 THE SYSTEM A*X = B IS PROBABLY NOT	209.
C		COMPATIBLE. A LEAST-SQUARES SOLUTION HAS	210.
C		BEEN OBTAINED WHICH IS AS ACCURATE AS SEEMS	211.
C		REASONABLE ON THIS MACHINE.	212.
C			213.
C		6 COND(ABAR) SEEMS TO BE SO LARGE THAT THERE IS	214.
C		NOT MUCH POINT IN DOING FURTHER ITERATIONS,	215.
C		GIVEN THE PRECISION OF THIS MACHINE.	216.
C		THERE COULD BE AN ERROR IN SUBROUTINE APROD.	217.
C			218.
C		7 THE ITERATION LIMIT ITNLIM WAS REACHED.	219.
C			220.
C	ANORM	OUTPUT AN ESTIMATE OF THE FROBENIUS NORM OF ABAR.	221.
C		THIS IS THE SQUARE-ROOT OF THE SUM OF SQUARES	222.
C		OF THE ELEMENTS OF ABAR.	223.
C		IF DAMP IS SMALL AND IF THE COLUMNS OF A	224.
C		HAVE ALL BEEN SCALED TO HAVE LENGTH 1.0,	225.
C		ANORM SHOULD INCREASE TO ROUGHLY SQRT(N).	226.

C		A RADICALLY DIFFERENT VALUE FOR ANORM MAY	227.
C		INDICATE AN ERROR IN SUBROUTINE APROD (THERE	228.
C		MAY BE AN INCONSISTENCY BETWEEN MODES 1 AND 2).	229.
C			230.
C	ACOND	OUTPUT AN ESTIMATE OF COND(ABAR), THE CONDITION	231.
C		NUMBER OF ABAR. A VERY HIGH VALUE OF ACOND	232.
C		MAY AGAIN INDICATE AN ERROR IN APROD.	233.
C			234.
C	RNORM	OUTPUT AN ESTIMATE OF THE FINAL VALUE OF NORM(RBAR),	235.
C		THE FUNCTION BEING MINIMIZED (SEE NOTATION	236.
C		ABOVE). THIS WILL BE SMALL IF $A \cdot X = B$ HAS	237.
C		A SOLUTION.	238.
C			239.
C	ARNORM	OUTPUT AN ESTIMATE OF THE FINAL VALUE OF	240.
C		NORM(ABAR(TRANPOSE)*RBAR), THE NORM OF	241.
C		THE RESIDUAL FOR THE USUAL NORMAL EQUATIONS.	242.
C		THIS SHOULD BE SMALL IN ALL CASES. (ARNORM	243.
C		WILL OFTEN BE SMALLER THAN THE TRUE VALUE	244.
C		COMPUTED FROM THE OUTPUT VECTOR X.)	245.
C			246.
C	XNORM	OUTPUT AN ESTIMATE OF THE NORM OF THE FINAL	247.
C		SOLUTION VECTOR X.	248.
C			249.
C			250.
C		SUBROUTINES AND FUNCTIONS USED	251.
C		-----	252.
C			253.
C	USER	APROD	254.
C	LSQR	NORMLZ	255.
C	BLAS	SCOPY, SNRM2, SSCAL (SEE LAWSON ET AL. BELOW)	256.
C		(SNRM2 IS USED ONLY IN NORMLZ)	257.
C	FORTRAN	ABS, MOD, SQRT	258.
C			259.
C			260.
C		PRECISION	261.
C		-----	262.
C			263.
C		THE NUMBER OF ITERATIONS REQUIRED BY LSQR WILL USUALLY DECREASE	264.
C		IF THE COMPUTATION IS PERFORMED IN HIGHER PRECISION. TO CONVERT	265.
C		LSQR AND NORMLZ BETWEEN SINGLE- AND DOUBLE-PRECISION, CHANGE	266.
C		THE WORDS	267.
C		SCOPY, SNRM2, SSCAL	268.
C		ABS, REAL, SQRT	269.
C		TO THE APPROPRIATE BLAS AND FORTRAN EQUIVALENTS.	270.
C			271.
C			272.
C		REFERENCES	273.
C		-----	274.
C			275.
C	PAIGE, C.C. AND SAUNDERS, M.A.	LSQR: AN ALGORITHM FOR SPARSE	276.
C		LINEAR EQUATIONS AND SPARSE LEAST SQUARES.	277.
C		ACM TRANSACTIONS ON MATHEMATICAL SOFTWARE 8, 1 (MARCH 1982).	278.
C			279.
C	LAWSON, C.L., HANSON, R.J., KINCAID, D.R. AND KROGH, F.T.		280.
C		BASIC LINEAR ALGEBRA SUBPROGRAMS FOR FORTRAN USAGE.	281.
C		ACM TRANSACTIONS ON MATHEMATICAL SOFTWARE 5, 3 (SEPT 1979),	282.
C		308-323 AND 324-325.	283.
C			284.
C			285.
C	LSQR.	THIS VERSION DATED 22 FEBRUARY 1982.	286.
C		-----	287.