BIT
1999, Vol. 39, No. 2, pp. 210–227

0006-3835/99/3902-0210 $15.00

BACKWARD ERROR BOUNDS FOR CONSTRAINED LEAST SQUARES PROBLEMS *

ANTHONY J. COX[1] and NICHOLAS J. HIGHAM[2†]

[1] *Present address: Ground Processing Department, Mail Point 31, Matra Marconi Space UK Anchorage Road, Portsmouth, PO3 5PU, UK. email: anthony.cox@mmsuk.co.uk*

[2] *Department of Mathematics, University of Manchester, Manchester, M13 9PL, England email: higham@ma.man.ac.uk. (http://www.ma.man.ac.uk/˜higham)*

Abstract.

We derive an upper bound on the normwise backward error of an approximate solution to the equality constrained least squares problem $\min_{Bx=d} \|b - Ax\|_2$. Instead of minimizing over the four perturbations to A, b, B and d, we fix those to B and d and minimize over the remaining two; we obtain an explicit solution of this simplified minimization problem. Our experiments show that backward error bounds of practical use are obtained when B and d are chosen as the optimal normwise relative backward perturbations to the constraint system, and we find that when the bounds are weak they can be improved by direct search optimization. We also derive upper and lower backward error bounds for the problem of least squares minimization over a sphere: $\min_{\|x\|_2 \le \alpha} \|b - Ax\|_2$.

AMS subject classification: 65F20

Key words: Equality constrained least squares problem, least squares minimization over a sphere, null space method, elimination method, method of weighting, backward error, backward stability.

1 Introduction.

Given an approximate solution to a mathematical problem, backward error analysis aims to answer the question “how much do we have to perturb the problem in order for the approximate solution to be an exact solution?” For a suitable measure of the size of the perturbation, the aim is to derive a formula for the backward error that is inexpensive to compute (or at least to bound) and is compatible with existing backward error analysis of algorithms for the problem. The backward error formula can then be used to test the stability of algorithms and the sharpness of backward error bounds obtained from rounding error analysis.

*Received April 1998. Revised December 1998. Communicated by Åke Björck.

†The work of the first author was supported by an Engineering and Physical Sciences Research Council Research Studentship. The work of the second author was supported by Engineering and Physical Sciences Research Council grant GR/L76532.

Existing work on backward errors covers a variety of problems, including linear systems [12, 21, 22], the unconstrained least squares problem [25, 30], underdetermined linear systems [26], and eigenvalue problems [13, 24] (for more references see [10] and [15]). Normwise and componentwise measures of the perturbations have been considered, as well as backward errors that preserve structure in the data.

In this work we are concerned with backward errors for two constrained least squares problems: the least squares problem with equality constraints,

$$\text{LSE}: \quad \min_{Bx=d} \|b - Ax\|_2, \quad A \in \mathbb{R}^{m\times n}, \quad B \in \mathbb{R}^{p\times n}, \quad m+p \ge n \ge p, \tag{1.1}$$

and least squares minimization over a sphere,

$$\text{LSS}: \quad \min_{\|x\|_2 \le \alpha} \|b - Ax\|_2, \quad A \in \mathbb{R}^{m\times n}, \quad m \ge n, \quad \alpha \ge 0. \tag{1.2}$$

Computation of backward errors has not previously been considered for these problems. Yet various solution methods with different numerical stability properties exist and so it is desirable to be able to test the stability of given approximate solutions. One application is to testing software [16]. For example, the LAPACK testing routines evaluate backward errors where possible. In LAPACK 2.0, the driver routine `xgglse.f` for the LSE problem is tested only on problems constructed to have a zero residual, for which scaled norms of $b - Ax$ and $d - Bx$ are compared with a tolerance. It is clearly desirable to have suitable tests available for least squares problems with nonzero residuals, particularly during the development of software.

An explicit formula is known for the backward error of an approximate solution to the unconstrained least squares (LS) problem. Determining such a formula in the equality constrained case is a more difficult nonlinear optimization problem and we make some progress towards its solution by developing a computable upper bound for the backward error. Instead of optimizing over all four perturbation matrices and vectors we fix the perturbations to the constraints and then optimize over the perturbations to the linear system itself. For the LSS problem we characterize the backward error and show how to compute upper and lower bounds for it.

In Section 2 we summarize normwise backward error results for a linear system and for the LS problem. In Section 3 we give a result that enables us to parametrize the normal equations of the LSE and LSS problems. This parametrization is used in Section 4 to derive an upper bound for the backward error for the LSE problem and in Section 6 to obtain upper and lower bounds for the backward error for the LSS problem. We present numerical experiments in Section 5 to show that our upper bound for the LSE problem is of practical use and we describe an expensive but effective way of attempting to reduce the size of the bound.

2 Backward error for linear systems and the LS problem.

We begin by summarizing normwise backward error results for a linear system and for the LS problem $\min_x \|b - Ax\|_2$. In the first result, $\|\cdot\|$ denotes any vector norm and the corresponding subordinate matrix norm. The vector z is defined to be dual to y if $z^T y = \|x\|_D \|y\| = 1$, where $\|z\|_D = \max_{x\neq 0} |z^T x|/\|x\|$.

THEOREM 2.1. (Rigal and Gaches [22]) *The normwise relative backward error*

$$\tau(y) := \min\{\, \epsilon : (A+E)y = b+f, \quad \|E\| \le \epsilon\|A\|, \quad \|f\| \le \epsilon\|b\| \,\} \tag{2.1}$$

is given by

$$\tau(y) = \frac{\|r\|}{\|A\|\,\|y\| + \|f\|} \tag{2.2}$$

where $r = b - Ay$. Optimal perturbations are

$$E = \frac{\|A\|\,\|y\|}{\|A\|\,\|y\| + \|b\|} r z^T, \qquad f = -\frac{\|b\|}{\|A\|\,\|y\| + \|b\|} r,$$

where z is a vector dual to y.

An explicit formula for a normwise backward error for the LS problem was discovered only recently. A superscript '+' denotes the pseudo-inverse; in particular, for a nonzero vector x, $x^+ = x^T/(x^T x)$.

THEOREM 2.2. (Waldén, Karlson and Sun [30]) *Let $A \in \mathbb{R}^{m\times n}$ $(m \ge n)$, $b \in \mathbb{R}^m$, $0 \ne y \in \mathbb{R}^n$, and $r = b - Ay$. The normwise backward error*

$$\eta(y) := \min\{\, \|[E \;\, \theta f]\|_F : \|(A+E)y - (b+f)\|_2 = \min \,\} \tag{2.3}$$

is given by

$$\eta(y) = \begin{cases} \dfrac{\|r\|_2}{\|y\|_2}\sqrt{\mu}, & \lambda_* \ge 0, \\[2ex] \left(\dfrac{\|r\|_2^2}{\|y\|_2^2}\mu + \lambda_* \right)^{1/2}, & \lambda_* < 0, \end{cases} \tag{2.4}$$

where

$$\lambda_* = \lambda_{\min}\left(AA^T - \mu\frac{rr^T}{\|y\|_2^2} \right), \quad \mu = \frac{\theta^2\|y\|_2^2}{1+\theta^2\|y\|_2^2}.$$

The corresponding perturbations of A and b are

$$(E_*, f_*) = \begin{cases} \left(\mu r y^+, -\dfrac{r}{1+\theta^2\|y\|_2^2} \right), & \lambda_* \ge 0, \\[2ex] \left(\mu r y^+ - vv^+(A + \mu r y^+), -\dfrac{(I - vv^+)r}{1+\theta^2\|y\|_2^2} \right), & \lambda_* < 0, \end{cases} \tag{2.5}$$

where v is an eigenvector of the matrix $AA^T - \mu rr^T/\|y\|_2^2$ that corresponds to its smallest eigenvalue λ_*.

The parameter θ in (2.3) allows some freedom in the definition of backward error. In the limit $\theta \to \infty$ perturbations to A only are allowed, while $\theta = \|A\|_F/\|b\|_2$ produces a relative backward error.

Although elegant, the formula (2.4) is unsuitable for practical computation as it expresses the backward error as the square root of another quantity; that quantity will be computed with absolute error proportional to the unit roundoff and so the computed backward error will be of order at least the square root of the unit roundoff. Using the relation

$$\frac{\|r\|_2^2}{\|y\|_2^2}\mu + \lambda_{\min}\left(AA^T - \mu\frac{rr^T}{\|y\|_2^2}\right) = \sigma_{\min}([A\ R])^2, \quad R := \sqrt{\mu}\frac{\|r\|_2}{\|y\|_2}(I - rr^+), \tag{2.6}$$

where $\sigma_{\min}$ denotes the smallest singular value, the formula can be rewritten as follows.

COROLLARY 2.3. (Waldén, Karlson and Sun [30]) *With the notation of Theorem* 2.2*, we have*

$$\eta(y) = \min\{\, \phi,\, \sigma_{\min}\big([\,A \quad \phi(I - rr^+)\,]\big)\,\}, \qquad \phi = \sqrt{\mu}\frac{\|r\|_2}{\|y\|_2}, \tag{2.7}$$

and the optimal perturbations E_* and f_* are given by (2.5) *with v a left singular vector of $[\,A \quad \phi(I - rr^+)\,]$ corresponding to the smallest singular value.*

As noted by Gu [11], the formula (2.7) has the disadvantage that the matrix $[\,A \quad \phi(I - rr^+)\,]$ has badly scaled columns when ϕ is large, which can cause standard methods for computing the singular value decomposition (SVD) to return an inaccurate computed smallest singular value. However, SVD methods developed in [8, 17] do not suffer from this difficulty as they are insensitive to the column scaling.

The formula (2.7) requires computation of the SVD of an $m \times (n+m)$ matrix, which can be prohibitively expensive for large problems. Recent attention has been focussed on obtaining more easily computable bounds on the backward error. Karlson and Waldén [18] derive upper and lower bounds for the quantity $\eta(y)|_{\theta=\infty}$ in which only A is perturbed. Their bounds can be computed in $O(mn)$ operations but the bounds can differ by an arbitrary factor. Gu [11] derives an approximation to $\eta(y)|_{\theta=\infty}$ that differs from it by a factor less than 2 and can be computed in $O(mn^2)$ operations.

3 Parametrizing the normal equations.

The proof of Theorem 2.2 given in [30] is based on a parametrization of the set of perturbations E that satisfy the normal equations $(A+E)^T(b-(A+E)y) = 0$. In this section we obtain a parametrization of a more general set of equations that forms part of the normal equations for the LSE problem (see Lemma 4.1) and the LSS problem (see (6.3) and (6.4)).

LEMMA 3.1. *Let $A \in \mathbb{R}^{m\times n}$, $C \in \mathbb{R}^{n\times p}$ and $y \in \mathbb{R}^n$, and define the sets*

$$\mathcal{E}(y) = \{\, E : (A+E)^T(b - (A+E)y) = Cu \text{ for some } u \in \mathbb{R}^p \,\}$$

and

$$\overline{\mathcal{E}}(y) = \{\, v(w^T C^T - v^+ A) + (I - vv^+)(ry^+ + Z(I - yy^+)) : \\ v \in \mathbb{R}^m,\ w \in \mathbb{R}^p,\ Z \in \mathbb{R}^{m\times n}\,\},$$

where $r = b - Ay$. *Then* $\mathcal{E}(y) = \overline{\mathcal{E}}(y)$.

PROOF. We begin by noting the identity, for any y and v,

$$E = (I - vv^+)Eyy^+ + vv^+E + (I - vv^+)E(I - yy^+). \tag{3.1}$$

We suppose that $E \in \mathcal{E}(y)$ and show how this expression for E simplifies. We have

$$Cu = (A + E)^T(b - (A + E)y) = (A + E)^T(r - Ey).$$

Defining

$$v = r - Ey, \tag{3.2}$$

we have $v^TA + v^TE = u^TC^T$, which can be written

$$v^+A + v^+E = w^TC^T, \qquad w = \begin{cases} u, & v = 0, \\ u/(v^Tv), & v \neq 0. \end{cases} \tag{3.3}$$

Premultiplying (3.2) by $I - vv^+$ gives

$$(I - vv^+)Ey = (I - vv^+)r. \tag{3.4}$$

Combining (3.1), (3.3) and (3.4) we obtain

$$\begin{aligned} E &= (I - vv^+)ry^+ - vv^+A + vw^TC^T + (I - vv^+)E(I - yy^+) \\ &= v(w^TC^T - v^+A) + (I - vv^+)(ry^+ + Z(I - yy^+)), \end{aligned}$$

where we have replaced E on the right-hand side by the arbitrary matrix Z. Hence $E \in \overline{\mathcal{E}}(y)$.

Conversely, let $E \in \overline{\mathcal{E}}(y)$. Then

$$\begin{aligned} Ey &= v(w^TC^T - v^+A)y + (I - vv^+)r \\ &= v(w^TC^Ty) - vv^+(Ay + r) + r \\ &= v(w^TC^Ty) - vv^+b + r. \end{aligned} \tag{3.5}$$

Furthermore, $E^Tv = v^Tv(Cw - A^Tv^{+T})$, so that $(A+E)^Tv = (v^Tv)Cw$. Hence, using (3.5),

$$\begin{aligned} (A + E)^T(b - (A + E)y) &= (A + E)^T(r - Ey) \\ &= (A + E)^T(-v(w^TC^Ty) + vv^+b) \\ &= (-w^TC^Ty + v^+b)(v^Tv)Cw \\ &=: Cz, \end{aligned}$$

so that $E \in \mathcal{E}(y)$. □

Lemma 3.1 shows that a general element of $\mathcal{E}(y)$ has the form

$$\begin{aligned} E &= v(w^TC^T - v^+A) + (I - vv^+)(ry^+ + Z(I - yy^+)) \\ &=: E_1 + E_2. \end{aligned} \tag{3.6}$$

Our next result gives the Frobenius norm of E.

LEMMA 3.2. *For E in* (3.6) *we have*

$$\begin{aligned}\|E\|_F^2 &= \|E_1\|_F^2 + \|E_2\|_F^2,\\ \|E_1\|_F^2 &= \|v\|_2^2 \; \|w^T C^T - v^+ A\|_2^2,\\ \|E_2\|_F^2 &= \frac{\|r\|_2^2 - v^+ r r^T v}{\|y\|_2^2} + \operatorname{trace}\big(Z(I - yy^+)Z^T\big) - v^+ Z(I - yy^+)Z^T v.\end{aligned}$$

The minimum value of $\|E\|_F$ over all Z is achieved when $Z = 0$.

PROOF. The proof is very similar to analysis in [30]. The norm equalities are a straightforward computation using the relations $\|E\|_F^2 = \operatorname{trace}(E^T E)$, $E_1^T E_2 = 0$ and $\operatorname{trace}(AB) = \operatorname{trace}(BA)$. For the last part, denote the eigenvalues of the positive semidefinite matrix $Z(I - yy^+)Z^T$ by $\lambda_1 \geq \cdots \geq \lambda_m$. Then

$$\operatorname{trace}(Z(I - yy^+)Z^T) = \sum_{i=1}^{m} \lambda_i \geq 0$$

and $\lambda_1 \geq v^+ Z(I - yy^+)Z^T v \geq 0$. Hence

$$\operatorname{trace}(Z(I - yy^+)Z^T) - v^+ Z(I - yy^+)Z^T v \geq \sum_{i=2}^{m} \lambda_i \geq 0,$$

with equality for $Z = 0$, as required. □

4 The LSE problem.

First, we state the normal equations for the LSE problem, which we recall is

$$\text{(4.1)} \qquad \min_{Bx=d} \|b - Ax\|_2, \qquad A \in \mathbb{R}^{m \times n}, \quad B \in \mathbb{R}^{p \times n}, \quad m + p \geq n \geq p.$$

LEMMA 4.1. *The vector x solves the LSE problem* (4.1) *if and only if it satisfies*

$$\begin{aligned} A^T(b - Ax) &= B^T \xi,\\ Bx &= d,\end{aligned}$$

where $\xi \in \mathbb{R}^p$ is a vector of Lagrange multipliers.

PROOF. This is a standard result that can be derived by considering the first and second order optimality conditions in the theory of constrained optimization. □

Let y be an approximate solution to the LSE problem (4.1). We define the set of quadruples of perturbations

$$\text{(4.2)} \quad \mathcal{E}_{\mathrm{LSE}}(y) = \left\{ (E, f, F, g) : y \text{ solves } \min_{(B+F)x = d+g} \|b + f - (A + E)x\|_2 \right\}.$$

A natural way to define the backward error of y is to minimize some norm $\|\cdot\|_W$ of the matrix $\left[\begin{smallmatrix} E & f \\ F & g \end{smallmatrix}\right]$ over the set $\mathcal{E}_{\rm LSE}(y)$:

$$\beta(y) = \min\left\{ \left\| \begin{bmatrix} E & f \\ F & g \end{bmatrix} \right\|_W : (E,f,F,g) \in \mathcal{E}_{\rm LSE}(y) \right\}. \tag{4.3}$$

Unfortunately, this optimization problem appears to be substantially more difficult than in the unconstrained case and we do not know how to solve it explicitly. Therefore we take a simplified approach.

We assume that we have chosen perturbations F and g such that $(B+F)y = d+g$. *Taking F and g as fixed*, we then minimize $\|[E\ \theta f]\|_F$ subject to y solving the LSE problem

$$\min_{(B+F)x=d+g} \|b+f-(A+E)x\|_2.$$

In other words, we are computing

$$\rho(y) = \min\{\, \|\,[E \quad \theta f\,]\,\|_F : (E,f,F,g) \in \mathcal{E}_{\rm LSE}(y)\,\}. \tag{4.4}$$

Denoting the optimal perturbations by E_* and f_*, we clearly have the upper bound for the backward error

$$\beta(y) \le \left\| \begin{bmatrix} E_* & f_* \\ F & g \end{bmatrix} \right\|_W .$$

If we do not wish to perturb the constraints then we can set $F=0$, $g=0$ and, in order for $(E,f,F,g) \in \mathcal{E}_{\rm LSE}(y)$ to be possible, imagine redefining $d := By$; the minimization in (4.4) can then be expressed as

$$\rho(y) = \min\left\{ \|\,[E \quad \theta f\,]\,\|_F : \|b+f-(A+E)y\|_2 = \min_{Bx=By} \|b+f-(A+E)x\|_2 \right\}.$$

We now derive an explicit formula for $\rho(y)$. We introduce the notation $P_{N(A)} = I - A^+A$, which is the projector onto the null space of A.

THEOREM 4.2. *For the LSE problem* (4.1) *let* $0 \ne y \in \mathbb{R}^n$ *and* $r = b - Ay$ *and assume that* F *and* g *satisfy* $(B+F)y = d+g$. *Then*

$$\rho(y) = \begin{cases} \dfrac{\|r\|_2}{\|y\|_2}\sqrt{\mu}, & \lambda_* \ge 0, \\[2ex] \left(\dfrac{\|r\|_2^2}{\|y\|_2^2}\mu + \lambda_* \right)^{1/2}, & \lambda_* < 0, \end{cases}$$

where

$$\lambda_* = \lambda_{\min}\left(A P_{N(B+F)} A^T - \mu \frac{rr^T}{\|y\|_2^2} \right), \qquad \mu = \frac{\theta^2 \|y\|_2^2}{1+\theta^2\|y\|_2^2}.$$

The corresponding perturbations of A and b are

$$(E_*, f_*) = \begin{cases} \left(\mu r y^+, -\dfrac{r}{1+\theta^2\|y\|_2^2} \right), & \lambda_* \ge 0, \\[2ex] \left(\mu r y^+ - vv^+(AP_{N(B+F)} + \mu r y^+), -\dfrac{(I - vv^+)r}{1+\theta^2\|y\|_2^2} \right), & \lambda_* < 0, \end{cases} \tag{4.5}$$

where v is an eigenvector of the matrix $AP_{N(B+F)}A^T - \mu rr^T/\|y\|_2^2$ that corresponds to its smallest eigenvalue λ_*.

PROOF. We first solve the simpler problem in which perturbations are restricted to A: we find

$$\gamma(y) = \min\{\, \|E\|_F : (E, 0, F, g) \in \mathcal{E}_{\mathrm{LSE}}(y)\,\}.$$

From Lemma 4.1 we see that the feasible E are those satisfying

$$(A+E)^T(b-(A+E)y) = (B+F)^T\xi$$

for some ξ. Lemma 3.1 shows that

$$E = v(w^T(B+F) - v^+A) + (I - vv^+)(ry^+ + Z(I - yy^+)) =: E_1 + E_2,$$

where v, w and Z are arbitrary. Lemma 3.2 gives, on setting $Z = 0$,

$$\|E_1\|_F^2 = \|v\|_2^2\|w^T(B+F) - v^+A\|_2^2, \qquad \|E_2\|_F^2 = \frac{\|r\|_2^2 - v^+rr^Tv}{\|y\|_2^2}.$$

For given v, $\|E_1\|_F^2$ is minimized by $w^T = v^+A(B+F)^+$ and the minimal value is

$$\|E_1\|_F^2 = \|v\|_2^2\|v^+AP_{N(B+F)}\|_2^2.$$

The remaining calculation is

$$\begin{aligned}\min_v \|E\|_F^2 &= \min_v\left(\frac{\|r\|_2^2}{\|y\|_2^2} + v^+\left(AP_{N(B+F)}A^T - \frac{rr^T}{\|y\|_2^2}\right)v\right)\\ &= \frac{\|r\|_2^2}{\|y\|_2^2} + \min\left(\lambda_{\min}\left(AP_{N(B+F)}A^T - \frac{rr^T}{\|y\|_2^2}\right), 0\right).\end{aligned}$$

Hence

$$\gamma(y) = \begin{cases} \dfrac{\|r\|_2}{\|y\|_2}, & \lambda_{\min} \geq 0,\\[2ex] \left(\dfrac{\|r\|_2^2}{\|y\|_2^2} + \lambda_{\min}\right)^{1/2}, & \lambda_{\min} < 0,\end{cases}$$

where

$$\lambda_{\min} = \lambda_{\min}\left(AP_{N(B+F)}A^T - \frac{rr^T}{\|y\|_2^2}\right).$$

The minimizing perturbation $E_{\min}$ is given by

$$E_{\min} = \begin{cases} ry^+, & \lambda_{\min} \geq 0,\\[1ex] ry^+ - vv^+(AP_{N(B+F)} + ry^+), & \lambda_{\min} < 0,\end{cases}$$

where v is an eigenvector of $AP_{N(B+F)}A^T - rr^T/\|y\|_2^2$ corresponding to $\lambda_{\min}$.

Finally, we reason in an identical manner to the proof of Corollary 2.1 in [30] to incorporate perturbations to b and thus to obtain the required result. $\square$

Our formulae for $\rho(y)$ and the corresponding optimal perturbations differ from those for $\eta(y)$ for the LS problem in Theorem 2.2 only by the presence of the term $P_{N(B+F)}$. It is therefore natural to ask whether Theorem 4.2 can be obtained directly from Theorem 2.2. Indeed, the LSE problem $\min_{Bx=d} \|b-Ax\|_2$ is equivalent to the LS problem $\min_x \|(b-AB^+d)-AP_{N(B)}x\|_2$. However, the coefficient matrix of this unconstrained problem is $AP_{N(B)}$, with $P_{N(B)}$ idempotent rather than orthogonal, and A appears in both the coefficient matrix and the right-hand side, so backward errors for this LS problem are not directly translatable into backward errors for the LSE problem.

Just as for the LS problem, we can rewrite the formula for $\rho(y)$ in a form better suited to finite precision computation.

COROLLARY 4.3. *With the notation of Theorem* 4.2, *we have*

$$\rho(y) = \min\{\, \phi,\, \sigma_{\min}\big([\, AP_{N(B+F)} \quad \phi(I-rr^+)\,]\big)\,\}, \qquad \phi = \sqrt{\mu}\frac{\|r\|_2}{\|y\|_2},$$

and the optimal perturbations E_* and f_* are given by (4.5) *with v a left singular vector of $[\, AP_{N(B+F)} \quad \phi(I-rr^+)\,]$ corresponding to the smallest singular value.*

Since the formula for the backward error bound involves the projector $P_{N(B+F)}$, there are potential difficulties in evaluating the bound accurately when $B+F$ is ill conditioned. We investigate this issue further in the next section.

Now we consider the choice of F and g. Of course, minimizing some suitable measure of the size of the triple $(F, g, \rho(y))$ over all F and g would give a true backward error, but this minimization is intractable. Natural choices for F and g are the optimal normwise relative backward perturbations to the constraint system, using the 2-norm. From Theorem 2.1 we have

$$F_* = \frac{\|B\|_2\|y\|_2}{\|B\|_2\|y\|_2+\|d\|_2} r_B y^+, \qquad g_* = \frac{-\|d\|_2}{\|B\|_2\|y\|_2+\|d\|_2} r_B, \tag{4.6}$$

where $r_B = d - By$, for which

$$\|F_*\|_2 = \tau\|B\|_2, \qquad \|g_*\|_2 = \tau\|d\|_2, \qquad \tau = \tau(y) = \frac{\|r_B\|_2}{\|B\|_2\|y\|_2+\|d\|_2}.$$

Returning to our definition (4.3) of backward error $\beta(y)$ for the LSE problem, let

$$\left\| \begin{bmatrix} E & f \\ F & g \end{bmatrix} \right\|_W = \max\left\{ \frac{\|E\|_2}{\|A\|_2}, \frac{\|f\|_2}{\|b\|_2}, \frac{\|F\|_2}{\|B\|_2}, \frac{\|g\|_2}{\|d\|_2} \right\}, \tag{4.7}$$

which makes $\beta(y)$ a normwise relative backward error. The following corollary gives our computable upper bound for $\beta(y)$.

COROLLARY 4.4. *For the norm* (4.7), *with F_* and g_* defined in* (4.6) *and E_* and f_* defined in terms of F_* and g_* by* (4.5),

$$\beta^U(y) := \left\| \begin{bmatrix} E_* & f_* \\ F_* & g_* \end{bmatrix} \right\|_W \geq \beta(y).$$

5 Numerical experiments.

We have carried out numerical experiments in an attempt to answer several questions arising from the analysis of the previous section.

1. Is the bound $\beta(y) \le \beta^U(y)$ sharp enough for $\beta^U(y)$ to confirm backward stability of an algorithm?
2. Is the bound $\beta^U(y)$ able to distinguish between stable and unstable algorithms for solving the LSE problem?
3. In cases where $\beta^U(y)$ is a weak bound, can we obtain a smaller upper bound for $\beta(y)$?
4. How sensitive is the evaluation of $\beta^U(y)$ to rounding errors?

Our experiments were performed in MATLAB, using simulated single precision arithmetic in which the result of every arithmetic operation was rounded to 24 bits; thus the unit roundoff $u = 2^{-24} \approx 5.96 \times 10^{-8}$.

We report results for LSE problems with $m = 10$, $n = 7$ and $p = 3$, with A and B random matrices with preassigned 2-norm condition numbers and geometrically distributed singular values, generated by the routine `randsvd`[1] from the Test Matrix Toolbox [14]. However, in one of the tests B was formed in this way with $\kappa_2(B) = 10$ and then $B(1\colon p, 1\colon p)$ was replaced by a matrix of random numbers from the normal(0,1) distribution all multiplied by 10^{-8}; this choice of B is denoted "small $\sigma_{\min}(B_{11})$". The vectors b and d have random elements from the normal(0,1) distribution. We solved the LSE problems using three different methods. Computed solutions are denoted by $\widehat{x}$.

1. The elimination method of Björck and Golub [4], [3, Section 5.1.2], which uses a QR factorization of B to eliminate p of the unknowns, thus reducing the problem to an unconstrained LS problem; backward error analysis for the method is given in [7]. It is well known that column pivoting needs to be used in the QR factorization in order to obtain a stable method (in fact, the method can break down otherwise). We tried the method both with and without column pivoting, to see whether the backward error bound $\beta^U(\widehat{x})$ could distinguish the difference in stability.
2. The null space method, as implemented in LAPACK using generalized QR factorization [1]. Backward error analysis of the null space method is given in [6].
3. The method of weighting, which approximates the solution to the LSE problem by the solution of the unconstrained LS problem

$$\min_x \left\| \begin{bmatrix} \mu d \\ b \end{bmatrix} - \begin{bmatrix} \mu B \\ A \end{bmatrix} x \right\|_2 \qquad (5.1)$$

[1] This routine is also accessible as `gallery('randsvd',...)` in MATLAB 5.

for a suitably large positive weight μ [3, Section 5.1.5], [19, Chapter 22], [29]. The LS problem is solved by Householder QR factorization with column pivoting. We tried three choices of μ: $u^{-1/3}$, which is recommended by Barlow [2] in connection with a subsequent deferred correction iteration; $u^{-1/2}$, which is commonly recommended; and the much larger value $u^{-3/2}$ recommended by Stewart[2] [23]. Note that there are two sources of error with the method of weighting: rounding errors and truncation errors stemming from a finite choice of μ.

We evaluated the backward error bound $\beta^U(\widehat{x})$ in both single and double precision, with $\theta = \|A\|_F/\|b\|_2$, but we report only the values computed in double precision. The two computed values differed by up to a factor 10 and we have not observed substantially larger differences in further tests. Since it is usually only the order of magnitude of the backward error that is of interest, we conclude that there is reasonable justification for evaluating $\beta^U(\widehat{x})$ at the working precision, especially since the single precision values almost always err on the side of being too large, so that they are still upper bounds.

As well as the backward error bound we evaluated the forward error

$$e(x) = \frac{\|x - \widehat{x}\|_2}{\|x\|_2},$$

together with a quantity ϕ such that

$$e(x) \le \beta(\widehat{x})\phi + O(\beta(\widehat{x})^2),$$

for the backward error β given by (4.3) and (4.7). This is a perturbation bound from [6]; ϕ is an upper bound on a normwise relative condition number. The significance of this forward error bound is that it provides the lower bound

$$\beta(\widehat{x}) \gtrsim \frac{e(x)}{\phi}, \tag{5.2}$$

which can be used to help judge the quality of our upper bound $\beta^U(\widehat{x})$.

The results are presented in Tables 5.1–5.3. For the ill-conditioned problem in Table 5.1, $\beta^U(\widehat{x})$ shows all the methods to be backward stable. For the problem in Table 5.2 in which B has a leading principal submatrix with a small singular value, the value $\beta^U(\widehat{x}) = 3.6\times10^{-1}$ suggests that the elimination method without column pivoting is behaving unstably; confirmation that $\beta^U(\widehat{x})$ is not pessimistic is obtained from (5.2), which shows that indeed $\beta(\widehat{x}) \ge 3.2 \times 10^{-1}$.

For the ill-conditioned problem in Table 5.3 the bounds $\beta^U(\widehat{x})$ are all large and we suspected that some of them are pessimistic. Therefore we used direct search optimization to try to choose F and g to minimize $\beta^U(\widehat{x})$, with starting guesses F_* and g_* from (4.6). An implementation from the Test Matrix Toolbox [14] of the multi-directional search method of Dennis and Torczon [27, 28] was used.

[2] Note that the term ϵ_M^{-1} in (3.8) of [23] should be $\epsilon_M^{-3/2}$.

Table 5.1: Results for $\kappa(A) = 10^6$, $\kappa(B) = 10^6$, $\phi = 3.8e6$.

	$\beta^U(\widehat{x})$	$e(\widehat{x})$
Elimination (no column pivoting)	7.4e−8	1.0e−3
Elimination (with column pivoting)	6.4e−8	4.4e−2
Null space method	2.9e−8	2.1e−3
Method of weighting, $\mu = u^{-1/3}$	9.6e−7	2.8e−1
Method of weighting, $\mu = u^{-1/2}$	8.1e−8	4.5e−2
Method of weighting, $\mu = u^{-3/2}$	1.1e−7	4.4e−2

Table 5.2: Results for $\kappa(A) = 10$, small $\sigma_{\min}(B_{11})$, $\phi = 4.1e1$.

	$\beta^U(\widehat{x})$	$e(\widehat{x})$
Elimination (no column pivoting)	3.6e−1	1.3e0
Elimination (with column pivoting)	3.5e−8	4.0e−7
Null space method	4.2e−8	2.0e−7
Method of weighting, $\mu = u^{-1/3}$	4.9e−6	9.6e−5
Method of weighting, $\mu = u^{-1/2}$	4.7e−8	6.6e−7
Method of weighting, $\mu = u^{-3/2}$	3.1e−8	2.3e−7

The improved values of β^U are shown in parentheses. All are of order u except for the method of weighting with $u^{-1/3}$ and $u^{-1/2}$. The lower bound (5.2) yields no new information about $\beta(\widehat{x})$ in these two cases, so although we may suspect it in view of the forward errors, we cannot conclude that the method of weighting yields a large backward error for the two smaller values of μ.

We conclude that the answer to the first three questions posed at the start of this section is "yes"—the backward error bound β^U is indeed a useful tool for understanding the behaviour of algorithms for solving the LSE problem. Further evidence can be found in [7], where the bound is used to test the behaviour of elimination methods, with particular reference to problems with bad row scaling. In [7] the perturbations E_*, f_*, F_* and g_* are found to yield useful upper bounds on both normwise and row-wise backward errors.

Table 5.3: Results for $\kappa(A) = 10$, $\kappa(B) = 10^7$, $\phi = 2.7e7$.

	$\beta^U(\widehat{x})$	$e(\widehat{x})$
Elimination (no column pivoting)	1.4e−2 (9.0e−8)	8.0e−2
Elimination (with column pivoting)	1.6e−3 (3.5e−8)	1.2e−2
Null space method	1.3e−2 (1.3e−7)	1.4e−1
Method of weighting, $\mu = u^{-1/3}$	2.9e−3 (2.9e−3)	1.0e0
Method of weighting, $\mu = u^{-1/2}$	2.6e−4 (6.7e−5)	1.0e0
Method of weighting, $\mu = u^{-3/2}$	1.6e−3 (5.9e−8)	1.2e−2

6 The LSS problem.

The LSS problem

$$\min_{\|x\|_2 \le \alpha} \|b - Ax\|_2, \quad A \in \mathbb{R}^{m\times n}, \quad m \ge n, \quad \alpha \ge 0 \tag{6.1}$$

is a special case of least squares minimization subject to a quadratic inequality constraint

$$\text{LSQI}: \quad \min_{\|d - Bx\|_2 \le \alpha} \|b - Ax\|_2, \qquad A \in \mathbb{R}^{m\times n}, \quad B \in \mathbb{R}^{p\times n}, \qquad m + p \ge n \ge p, \quad \alpha \ge 0. \tag{6.2}$$

An important application of the LSS problem is in trust region methods in nonlinear optimization [20]. A comprehensive treatment of the LSS and LSQI problems is given by Gander [9] (see also Björck [3, Section 5.3] and Golub and Van Loan [10, Section 12.1]).

The minimum norm solution to the problem (6.1) without the constraint has norm

$$\gamma = \|A^+ b\|_2,$$

and if $\gamma \le \alpha$ then clearly this vector solves the LSS problem. If $\gamma > \alpha$ then the constraint $\|x\|_2 \le \alpha$ is active at the solution. Using standard theory of constrained optimization it is straightforward to show that x solves the LSS problem if and only if it satisfies the normal equations

$$\left.\begin{aligned} A^T(b - Ax) &= 0 \\ \|x\|_2 &\le \alpha \end{aligned}\right\} \quad \text{if} \quad \gamma \le \alpha, \tag{6.3}$$

$$\left.\begin{aligned} A^T(b - Ax) &= \xi x \\ \|x\|_2 &= \alpha \end{aligned}\right\} \quad \text{if} \quad \gamma > \alpha, \tag{6.4}$$

where ξ is a *nonnegative* Lagrange multiplier. Thus an LSS solution always exists, and it is unique if $\text{rank}(A) = n$.

Let y be an approximate LSS solution and define the set of triples of perturbations

$$\mathcal{E}_{\text{LSS}}(y) = \left\{ (E, f, \delta) : y \text{ solves } \min_{\|x\|_2 \le \alpha + \delta} \|b + f - (A + E)x\|_2 \right\}.$$

We define the backward error by

$$\beta(y) = \min\{ \|\, [E \quad \theta f \quad \phi\,\delta e_1]\,\|_F : (E, f, \delta) \in \mathcal{E}_{\text{LSS}}(y) \}, \tag{6.5}$$

where θ and ϕ are parameters and e_1 denotes the first unit vector. We also define

$$\psi_+(y) := \min\{ \|[E\ \theta f]\|_F : (A + E)^T(b + f - (A + E)y) = \xi y, \text{ for some } \xi \ge 0 \}, \tag{6.6}$$

$$\psi_0(y) := \min\{ \|[E\ \theta f]\|_F : (A + E)^T(b + f - (A + E)y) = 0 \}, \tag{6.7}$$

for which, clearly, $\psi_+(y) \le \psi_0(y)$. The quantity $\psi_0(y)$ is just the backward error for the standard LS problem, given by Theorem 2.2.

The backward error $\beta(y)$ can be calculated as follows. We denote the optimal perturbations in (6.5) by E_*, f_* and δ_* and those in (6.6) and (6.7) by E_+, f_+ and E_0, f_0, respectively. We write $\gamma_+ = \|(A+E_+)^+(b+f_+)\|_2$ and similarly for γ_0.

1. First, suppose that $\|y\|_2 < \alpha$. There are two possibilities and we take whichever yields the smaller value of $\|[E\ \theta f\ \phi\delta e_1]\|_F$. The first choice is $\delta = 0$, $E = E_0$ and $f = f_0$. Then y is the solution of the unconstrained problem for $A+E$ and $b+f$ and so $\gamma_0 \le \|y\|_2 < \alpha$, since γ_0 is the norm of the minimum norm solution; hence the normal equations (6.3) are satisfied. The second choice is $\delta = \|y\|_2 - \alpha < 0$, which puts y on the boundary of the perturbed constraint region, with $E = E_+$ and $f = f_+$; as long as $\gamma_+ > \alpha + \delta = \|y\|_2$, the normal equations (6.4) are satisfied and since $\psi_+(y) \le \psi_0(y)$ the bigger constraint perturbation could be counteracted by smaller perturbations to A and b.

2. Next, suppose that $\|y\|_2 = \alpha$. There is clearly no advantage to increasing α so we set $\delta_* = 0$. If $\gamma_+ > \alpha$ then $E_* = E_+$ and $f_* = f_+$, with the normal equations (6.4) being satisfied. Otherwise $E_* = E_0$ and $f_* = f_0$, with the relevant normal equations being (6.3).

3. The final case is $\|y\|_2 > \alpha$. We set $\delta_* = \|y\|_2 - \alpha$. If $\gamma_+ > \alpha + \delta_*$ then $E_* = E_+$ and $f_* = f_+$, and we are in case (6.4) of the normal equations. Otherwise $E_* = E_0$ and $f_* = f_0$, with the relevant normal equations being (6.3).

This procedure for obtaining the backward error requires us to determine the quantity $\psi_+(y)$. Unfortunately, we do not have a computable characterization of $\psi_+(y)$. However, we can compute the quantity

$$\psi(y) := \min\{\, \|[E\ \theta f]\|_F : (A+E)^T(b+f-(A+E)y) = \xi y, \text{ for some } \xi\,\}, \tag{6.8}$$

in view of the following theorem.

THEOREM 6.1. *For the LSS problem* (6.1) *let* $0 \ne y \in \mathbb{R}^n$ *and* $r = b - Ay$. *Then*

$$\psi(y) = \begin{cases} \dfrac{\|r\|_2}{\|y\|_2}\sqrt{\mu}, & \lambda_* \ge 0, \\[2ex] \left(\dfrac{\|r\|_2^2}{\|y\|_2^2}\mu + \lambda_*\right)^{1/2}, & \lambda_* < 0, \end{cases}$$

where

$$\lambda_* = \lambda_{\min}\left(A(I - yy^+)A^T - \mu\frac{rr^T}{\|y\|_2^2}\right), \qquad \mu = \frac{\theta^2\|y\|_2^2}{1+\theta^2\|y\|_2^2}.$$

The corresponding perturbations of A and b are

$$(6.9)\quad (E_*, f_*) = \begin{cases} \left(\mu r y^+, -\dfrac{r}{1+\theta^2\|y\|_2^2}\right), & \lambda_* \geq 0, \\ \left(\mu r y^+ - vv^+(A(I-yy^+)+\mu r y^+), -\dfrac{(I-vv^+)r}{1+\theta^2\|y\|_2^2}\right), & \lambda_* < 0, \end{cases}$$

where v is an eigenvector of the matrix $A(I-yy^+)A^T - \mu rr^T/\|y\|_2^2$ that corresponds to its smallest eigenvalue λ_*.

Proof. We first solve the simpler problem in which perturbations are restricted to A. In this case the feasible E satisfy

$$(A+E)^T(b-(A+E)y) = \xi y$$

for some ξ. Applying Lemma 3.1 with $p=1$ we find that

$$E = v(\omega y^+ - v^+A) + (I-vv^+)(ry^+ + Z(I-yy^+)) =: E_1 + E_2,$$

where v, ω and Z are arbitrary. Lemma 3.2 gives, on setting $Z=0$,

$$\|E_1\|_F^2 = \|v\|_2^2\|\omega y^+ - v^+A\|_2^2, \qquad \|E_2\|_F^2 = \frac{\|r\|_2^2 - v^+rr^Tv}{\|y\|_2^2}.$$

For given v, $\|E_1\|_F^2$ is minimized by $\omega = v^+Ay$ and the minimal value is

$$\|E_1\|_F^2 = \|v\|_2^2\|v^+A(I-yy^+)\|_2^2.$$

The rest of the proof proceeds in the same way as in the proof of Theorem 4.2. □

As for the LS and LSE problems, for numerical evaluation the formula for $\psi(y)$ in Theorem 6.1 should be rewritten in terms of an SVD. The relevant formula is (2.7) with A replaced by $A(I-yy^+)$.

We now have three quantities $\psi(y)$, $\psi_+(y)$ and $\psi_0(y)$ that satisfy

$$(6.10)\qquad \psi(y) \leq \psi_+(y) \leq \psi_0(y),$$

and so we can compute lower and upper bounds on $\psi_+(y)$ and hence on the backward error $\beta(y)$. If ξ is positive for the optimal E and f in (6.8) then $\psi(y) = \psi_+(y)$ and we obtain the backward error exactly. The exact solution x satisfies $\psi(x) = \psi_+(x) = 0$ and if the Lagrange multiplier corresponding to x is strictly positive (recall that it is always nonnegative) then there is a ball around x within which $\xi > 0$ and hence $\psi = \psi_+$. Therefore we are certainly able to compute $\beta(y)$ exactly for y sufficiently close to x, though we do not have a practically useful way to quantify "sufficiently close".

We give a numerical example to illustrate the analysis. Consider the problem from [10, Example 12.1.1] with

$$A = \begin{bmatrix} 2 & 0 \\ 0 & 1 \\ 0 & 0 \end{bmatrix}, \qquad b = \begin{bmatrix} 4 \\ 2 \\ 3 \end{bmatrix}, \qquad \alpha = 1.$$

The exact solution has 2-norm α and is given, to five significant figures, by $x = [0.93334 \;\; 0.35898]^T$, with Lagrange multiplier $\xi = 4.5713$. We consider three different approximate solutions, with parameters $\theta = 1$, $\phi = 1$ in (6.5).

1. $y = x - [10^{-10} \;\; 10^{-10}]^T$. Here, $\xi = 4.57 > 0$ at the optimum in (6.8), so $\psi(y) = \psi_+(y)$ and we can obtain the exact backward error. We have $\|y\|_2 < \alpha$ so we are in case 1 above. For the choice $\delta = \|y\|_2 - \alpha < 0$, $E = E_+$ and $f = f_+$ we have $\gamma_+ > \|y\|_2$ and this choice yields the smaller value of $\|[E \;\theta f \;\phi\delta e_1]\|_F$, giving $\beta(y) = 3.72 \times 10^{-10}$.

2. $y = [1 \;\; 1]^T$. Here, $\xi = 2.25 > 0$ at the optimum in (6.8), so again $\psi(y) = \psi_+(y)$ and we can obtain the exact backward error. We have $\|y\|_2 > \alpha$ and $\gamma_+ > \alpha + \delta_*$ in case 3 above, giving $\beta(y) = 1.01$.

3. $y = [-1 \;\; 1]^T$. Here, $\xi = -5.49 < 0$ at the optimum in (6.8), so $\psi(y) \neq \psi_+(y)$. We have $\|y\|_2 > \alpha$ and we can obtain bounds on $\beta(y)$ by considering case 3 above. A feasible set of perturbations is $\delta = \|y\|_2 - \alpha$, $E = E_0$ and $f = f_0$, which gives $\beta(y) \leq 2.36$. For $\delta = \|y\|_2 - \alpha$, and E and f the optimal perturbations in (6.8) we find $\|(A+E)^+(b+f)\|_2 > \alpha + \delta$, and so we have the lower bound $\beta(y) \geq 2.09$. Thus the backward error is known to within an order of magnitude.

7 Conclusions.

By extending analysis of Waldén, Karlson and Sun [30] we have obtained upper bounds on the normwise backward error of an arbitrary approximate solution to the LSE problem. Instead of minimizing over all four perturbations to A, b, B and d, we fix those to B and d and minimize over the remaining two. In our experiments we found that choosing B and d to be the optimal normwise relative backward perturbations often leads to a sharp backward error bound, and that when this bound is weak direct search minimization can reduce it significantly, albeit at considerable extra expense. Our approach has the advantage of providing useful information about the backward error where none was previously available, and it therefore can profitably be used in the development and analysis of algorithms (see [7] for an example).

We have also obtained upper and lower backward error bounds for the LSS problem. By combining our techniques for the LSE problem with those for the LSS problem we can obtain upper bounds for the backward error for the general LSQI problem (6.2) [5].

Evaluation of the backward error bounds is expensive. For the LSE problem, we have to compute a projector of a $p \times n$ matrix and the SVD of an $m \times (n+m)$ matrix. However, the techniques developed in [11] and [18] can be adapted to produce cheaper estimates of the bounds (for the LSE problem, simply replace A by $AP_{N(B+F)}$ in [11] and [18]).

REFERENCES

1. E. Anderson, Z. Bai, C. H. Bischof, J. W. Demmel, J. J. Dongarra, J. J. Du Croz, A. Greenbaum, S. J. Hammarling, A. McKenney, S. Ostrouchov, and D. C. Sorensen, *LAPACK Users' Guide, Release 2.0*, 2nd ed., Society for Industrial and Applied Mathematics, Philadelphia, PA, 1995.
2. J. L. Barlow, *Error analysis and implementation aspects of deferred correction for equality constrained least squares problems*, SIAM J. Numer. Anal., 25:6 (1988), pp. 1340–1358.
3. Å. Björck, *Numerical Methods for Least Squares Problems*, Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 1996.
4. Å. Björck and G. H. Golub, *Iterative refinement of linear least squares solutions by Householder transformation*, BIT, 7 (1967), pp. 322–337.
5. A. J. Cox, *Stability of Algorithms for Solving Weighted and Constrained Least Squares Problems*, PhD thesis, University of Manchester, Manchester, England, Oct. 1997.
6. A. J. Cox and N. J. Higham, *Accuracy and stability of the null space method for solving the equality constrained least squares problem*, BIT 39:1 (1999), pp. 34–50.
7. A. J. Cox and N. J. Higham, *Row-wise backward stable elimination methods for the equality constrained least squares problem*, Numerical Analysis Report No. 319, Manchester Centre for Computational Mathematics, Manchester, England, Mar. 1998. To appear in SIAM J. Matrix Anal. Appl.
8. J. W. Demmel, M. Gu, S. Eisenstat, I. Slapničar, K. Veselić, and Z. Drmač, *Computing the singular value decomposition with high relative accuracy*, Technical Report CS-97-348, Department of Computer Science, University of Tennessee, Knoxville, TN, Feb. 1997. LAPACK Working Note 119. Revised August 1998.
9. W. Gander, *Least squares with a quadratic constraint*, Numer. Math., 36 (1981), pp. 291–307.
10. G. H. Golub and C. F. Van Loan, *Matrix Computations*, 3rd ed., Johns Hopkins University Press, Baltimore, MD, 1996.
11. M. Gu, *Backward perturbation bounds for linear least squares problems*, SIAM J. Matrix Anal. Appl., 20:2 (1998), pp. 363–372.
12. D. J. Higham and N. J. Higham, *Backward error and condition of structured linear systems*, SIAM J. Matrix Anal. Appl., 13:1 (1992), pp. 162–175.
13. D. J. Higham and N. J. Higham, *Structured backward error and condition of generalized eigenvalue problems*, SIAM J. Matrix Anal. Appl., 20:2 (1998), pp. 493–512.
14. N. J. Higham, *The Test Matrix Toolbox for* MATLAB *(version 3.0)*, Numerical Analysis Report No. 276, Manchester Centre for Computational Mathematics, Manchester, England, Sept. 1995.
15. N. J. Higham, *Accuracy and Stability of Numerical Algorithms*, Society for Industrial and Applied Mathematics, Philadelphia, PA, 1996.
16. N. J. Higham, *Testing linear algebra software*, in Quality of Numerical Software: Assessment and Enhancement, R. F. Boisvert, ed., Chapman & Hall, London, 1997, pp. 109–122.
17. N. J. Higham, *QR factorization with complete pivoting and accurate computation of the SVD*, Numerical Analysis Report No. 324, Manchester Centre for Computational Mathematics, Manchester, England, Sept. 1998. Submitted to Linear Algebra and Appl.

18. R. Karlson and B. Waldén, *Estimation of optimal backward perturbation bounds for the linear least squares problem*, BIT, 37:4 (1997), pp. 862–869.
19. C. L. Lawson and R. J. Hanson, *Solving Least Squares Problems*, Society for Industrial and Applied Mathematics, Philadelphia, PA, 1995. Revised republication of work first published in 1974 by Prentice-Hall.
20. J. J. Moré and D. C. Sorensen, *Computing a trust region step*, SIAM J. Sci. Stat. Comput., 4:3 (1983), pp. 553–572.
21. W. Oettli and W. Prager, *Compatibility of approximate solution of linear equations with given error bounds for coefficients and right-hand sides*, Numer. Math., 6 (1964), pp. 405–409.
22. J. L. Rigal and J. Gaches, *On the compatibility of a given solution with the data of a linear system*, J. Assoc. Comput. Mach., 14:3 (1967), pp. 543–548.
23. G. W. Stewart, *On the weighting method for least squares problems with linear equality constraints*, BIT, 37:4 (1997), pp. 961–967.
24. J. Sun, *A note on backward perturbations for the Hermitian eigenvalue problem*, BIT, 35 (1995), pp. 385–393.
25. J. Sun, *Optimal backward perturbation bounds for the linear least-squares problem with multiple right-hand sides*, IMA J. Numer. Anal., 16:1 (1996), pp. 1–11.
26. J. Sun and Z. Sun, *Optimal backward perturbation bounds for underdetermined systems*, SIAM J. Matrix Anal. Appl., 18:2 (1997), pp. 393–402.
27. V. J. Torczon, *Multi-Directional Search: A Direct Search Algorithm for Parallel Machines*, PhD thesis, Rice University, Houston, TX, May 1989.
28. V. J. Torczon, *On the convergence of the multidirectional search algorithm*, SIAM J. Optimization, 1:1 (1991), pp. 123–145.
29. C. F. Van Loan, *On the method of weighting for equality-constrained least-squares problems*, SIAM J. Numer. Anal., 22:5 (1985), pp. 851–864.
30. B. Waldén, R. Karlson, and J. Sun, *Optimal backward perturbation bounds for the linear least squares problem*, Numer. Linear Algebra Appl., 2:3 (1995), pp. 271–286.