How to estimate the backward error in LSQR

Pavel Jiránek[1]

joint work with David Titley-Peloquin[2]

[1]CERFACS, Toulouse, France

[2]School of Computer Science, McGill University, Montreal, Canada

Joint French-Czech Workshop on

Krylov Methods for Inverse Problems

Institute of Computer Science
Prague, Czech Republic

July 19–20, 2010

Outline

Outline

Least squares problems

Given the matrix $A \in \mathbb{R}^{m \times n}$ with $\mathrm{rank}(A) = n$ and the vector $b \in \mathbb{R}^m$ the (linear) least squares problem is to find $\hat{x} \in \mathbb{R}^n$ such that

$$\|b - A\hat{x}\|_2 = \min_x \|b - Ax\|_2. \tag{LS}$$

The (unique) solution of (LS) can be found by solving the system of normal equations

$$A^T A x = A^T b \qquad \Rightarrow \qquad \hat{x} = (A^T A)^{-1} A^T b \equiv A^\dagger b.$$

Least squares problems

Given the matrix $A \in \mathbb{R}^{m \times n}$ with $\mathrm{rank}(A) = n$ and the vector $b \in \mathbb{R}^m$ the (linear) least squares problem is to find $\hat{x} \in \mathbb{R}^n$ such that

$$\|b - A\hat{x}\|_2 = \min_x \|b - Ax\|_2. \qquad \text{(LS)}$$

The (unique) solution of (LS) can be found by solving the system of normal equations

$$A^T A x = A^T b \qquad \Rightarrow \qquad \hat{x} = (A^T A)^{-1} A^T b \equiv A^\dagger b.$$

The least squares problem can be interpreted as the orthogonal decomposition of b onto $\mathcal{R}(A)$ and $\mathcal{N}(A^T) = \mathcal{R}(A)^\perp$:

$$b = P_A b + P_A^\perp b = A\hat{x} + \hat{r},$$

where $P_A \equiv AA^\dagger$ and $P_A^\perp = I - P_A$.

For useful background, see Björck [1996], [Golub and Van Loan, 1996, Chap. 5], Lawson and Hanson [1974].

Solution methods for sparse least squares problems

Solution methods for linear least squares problems are usually based on some algorithm applied (explicitly or implicitly) either on the symmetric positive definite system of normal equations

$$A^T A x = A^T b \qquad \Leftrightarrow \qquad A^T r = 0, \quad r = b - Ax,$$

or on the indefinite system

$$\begin{bmatrix} I & A \\ A^T & 0 \end{bmatrix} \begin{bmatrix} r \\ x \end{bmatrix} = \begin{bmatrix} b \\ 0 \end{bmatrix}.$$

Solution methods for sparse least squares problems

Solution methods for linear least squares problems are usually based on some algorithm applied (explicitly or implicitly) either on the symmetric positive definite system of normal equations

$$A^T A x = A^T b \qquad \Leftrightarrow \qquad A^T r = 0, \quad r = b - Ax,$$

or on the indefinite system

$$\begin{bmatrix} I & A \\ A^T & 0 \end{bmatrix} \begin{bmatrix} r \\ x \end{bmatrix} = \begin{bmatrix} b \\ 0 \end{bmatrix}.$$

Classification of solution methods for sparse problems:

- **Direct methods**: variants of the Gaussian elimination, QR factorization.
- **Iterative methods**: classical (stationary) methods, Krylov subspace methods.

See, e.g., Heath [1984], Björck [1996, Chap. 6 & 7].

Outline

The backward error for consistent problems I

We consider a system $Ax = b$ such that $b \in \mathcal{R}(A)$ and let $\tilde{x} \in \mathbb{R}^n$ be an approximation to its solution.

Let $\mathcal{E}_{\mathrm{CS}}(\tilde{x})$ be the set of all *backward perturbations* of the data in $Ax = b$ associated to $\tilde{x}$,

$$\mathcal{E}_{\mathrm{CS}}(\tilde{x}) \equiv \{[E, f];\ (A + E)\tilde{x} = b + f\}.$$

The backward error associated to $\tilde{x}$ is defined as a *"size"* of a minimal backward perturbation in $\mathcal{E}_{\mathrm{CS}}(\tilde{x})$.

The backward error for consistent problems II

Let α, β be positive and let

$$\xi_{\mathrm{CS}}(\tilde{x}, \alpha, \beta) \equiv \min\{\xi;\ [E, f] \in \mathcal{E}_{\mathrm{CS}}(\tilde{x}),\ \|E\|_F \leq \xi\alpha\|A\|_F,\ \|f\|_2 \leq \xi\beta\|b\|_2\}$$
$$= \min\left\{\left\|\left[\frac{\|E\|_F}{\alpha\|A\|_F}, \frac{\|f\|_2}{\beta\|b\|_2}\right]\right\|_\infty; [E, f] \in \mathcal{E}_{\mathrm{CS}}(\tilde{x})\right\}.$$

Then $\tilde{x}$ solves a nearby problem

$$(A + E)\tilde{x} = b + f, \qquad \|E\|_F \leq \alpha\|A\|_F, \quad \|f\|_2 \leq \beta\|f\|_2$$

if and only if $\xi_{\mathrm{CS}}(\tilde{x}) \leq 1$.

The backward error for consistent problems II

Let α, β be positive and let

$$\xi_{\mathrm{CS}}(\tilde{x}, \alpha, \beta) \equiv \min\{\xi;\ [E, f] \in \mathcal{E}_{\mathrm{CS}}(\tilde{x}),\ \|E\|_F \leq \xi\alpha\|A\|_F,\ \|f\|_2 \leq \xi\beta\|b\|_2\}$$
$$= \min\left\{\left\|\left[\frac{\|E\|_F}{\alpha\|A\|_F}, \frac{\|f\|_2}{\beta\|b\|_2}\right]\right\|_\infty ; [E, f] \in \mathcal{E}_{\mathrm{CS}}(\tilde{x})\right\}.$$

Then $\tilde{x}$ solves a nearby problem

$$(A + E)\tilde{x} = b + f, \qquad \|E\|_F \leq \alpha\|A\|_F, \quad \|f\|_2 \leq \beta\|f\|_2$$

if and only if $\xi_{\mathrm{CS}}(\tilde{x}) \leq 1$.

The problem was solved by Rigal and Gaches [1967] and the expression for $\xi_{\mathrm{CS}}(\tilde{x}; \alpha, \beta)$ is given by

$$\xi_{\mathrm{CS}}(\tilde{x}; \alpha, \beta) = \frac{\|\tilde{r}\|_2}{\alpha\|A\|_F\|\tilde{x}\|_2 + \beta\|b\|_2}, \qquad \tilde{r} \equiv b - A\tilde{x}.$$

If $\alpha = \mathcal{O}(\varepsilon)$, $\beta = \mathcal{O}(\varepsilon)$ the $\tilde{x}$ is usually called the *backward stable solution* of $Ax = b$ in finite precision arithmetic characterized by the machine precision ε.

The backward error for consistent problems III

We may also consider

$$\mu_{\mathrm{CS}}(\tilde{x};\theta) = \min\{\|[E,\theta f]\|_F;\ [E,f] \in \mathcal{E}_{\mathrm{CS}}(\tilde{x})\}$$

for some $\theta > 0$. The explicit formula for $\mu_{\mathrm{CS}}(\tilde{x};\theta)$ is

$$\mu_{\mathrm{CS}}(\tilde{x};\theta) = \frac{\theta\|\tilde{r}\|_2}{\sqrt{1+\theta^2\|\tilde{x}\|_2^2}}.$$

The backward error for consistent problems III

We may also consider

$$\mu_{\mathrm{CS}}(\tilde{x};\theta) = \min\{\|[E, \theta f]\|_F;\ [E, f] \in \mathcal{E}_{\mathrm{CS}}(\tilde{x})\}$$

for some $\theta > 0$. The explicit formula for $\mu_{\mathrm{CS}}(\tilde{x};\theta)$ is

$$\mu_{\mathrm{CS}}(\tilde{x};\theta) = \frac{\theta\|\tilde{r}\|_2}{\sqrt{1+\theta^2\|\tilde{x}\|_2^2}}.$$

Both ξ_{CS} and μ_{CS} are equivalent in the sense that

$$\frac{1}{\sqrt{2}}\frac{\mu_{\mathrm{CS}}(\tilde{x};\theta)}{\alpha\|A\|_F} \le \xi_{\mathrm{CS}}(\tilde{x};\alpha,\beta) \le \frac{\mu_{\mathrm{CS}}(\tilde{x};\theta)}{\alpha\|A\|_F}, \qquad \theta = \frac{\alpha\|A\|_F}{\beta\|b\|_2}.$$

The backward error for LS problems I

For LS problems, we have

$$\mathcal{E}_{\mathrm{LS}}(\tilde{x}) \equiv \{[E, f];\ (A+E)^T[(b+f)-(A+E)\tilde{x}]=0\}.$$

Obviously, $\mathcal{E}_{\mathrm{CS}}(\tilde{x}) \subset \mathcal{E}_{\mathrm{LS}}(\tilde{x})$.

We may define $\xi \equiv \xi_{\mathrm{LS}}$ and $\mu \equiv \mu_{\mathrm{LS}}$ in the similar way as for consistent problems.

Note that the same equivalence of ξ and μ holds as for consistent problems,

$$\frac{1}{\sqrt{2}}\frac{\mu(\tilde{x};\theta)}{\alpha\|A\|_F} \leq \xi(\tilde{x};\alpha,\beta) \leq \frac{\mu(\tilde{x};\theta)}{\alpha\|A\|_F}, \qquad \theta = \frac{\alpha\|A\|_F}{\beta\|b\|_2}.$$

The backward error for LS problems II

The solution to the optimal backward perturbation problem for LS problems was an open problem for a long time. It was shown by Stewart [1977] that

$$E_1 = -\tilde{r}\tilde{r}^\dagger A, \qquad E_2 = (\tilde{r} - \hat{r})\tilde{x}^\dagger = P_A\tilde{r}\tilde{x}^\dagger,$$

are backward perturbations of A corresponding to $\tilde{x}$, i.e.,

$$(A + E_i)^T[b - (A + E_i)\tilde{x}] = 0, \qquad i = 1, 2,$$

but none of them needs to be optimal. The norms of E_1 and E_2 are given by

$$\|E_1\|_2 = \|E_1\|_F = \frac{\|A^T\tilde{r}\|_2}{\|\tilde{r}\|_2}, \qquad \|E_2\|_2 = \|E_2\|_F = \frac{\|\hat{r} - \tilde{r}\|_2}{\|\tilde{x}\|_2} = \frac{\|P_A\tilde{r}\|_2}{\|\tilde{x}\|_2}.$$

The backward error for LS problems III

The expression for $\mu(\tilde{x};\theta)$ was provided by Waldén, Karlson, and Sun [1995].

Let

$$\tilde{r} \equiv b - A\tilde{x}, \qquad \omega \equiv \frac{\theta\|\tilde{r}\|_2}{\sqrt{1+\theta^2\|\tilde{x}\|_2^2}} = \mu_{\mathrm{CS}}(\tilde{x};\theta), \qquad N \equiv [A, \omega(I - \tilde{r}\tilde{r}^\dagger)].$$

Then

$$\mu(\tilde{x};\theta) = \min\{\omega, \sigma_{\min}(N)\}.$$

If $b \notin \mathcal{R}(A)$, then $\mu(\tilde{x};\theta) < \mu_{\mathrm{CS}}(\tilde{x};\theta)$, i.e., $\mu(\tilde{x};\theta) = \sigma_{\min}(N)$.

The backward error for LS problems IV

In comparison with consistent problems, the backward error for LS problems is much harder to compute: it requires to compute the smallest singular value of the $m \times (m+n)$ matrix N.

Some attempts to derive tight lower and upper bounds for μ which are cheaper to compute appeared in Waldén, Karlson, and Sun [1995], Karlson and Waldén [1997].

The backward error for LS problems IV

In comparison with consistent problems, the backward error for LS problems is much harder to compute: it requires to compute the smallest singular value of the $m \times (m+n)$ matrix N.

Some attempts to derive tight lower and upper bounds for μ which are cheaper to compute appeared in Waldén, Karlson, and Sun [1995], Karlson and Waldén [1997].

The estimate

$$\nu(x;\theta) \equiv \frac{\omega}{\|\tilde{r}\|_2} \left\| \begin{bmatrix} A \\ \omega I \end{bmatrix} \begin{bmatrix} A \\ \omega I \end{bmatrix}^{\dagger} \begin{bmatrix} \tilde{r} \\ 0 \end{bmatrix} \right\|_2$$

was studied by Gu [1998] (see also Karlson and Waldén [1997]) for which

$$\frac{\|\hat{r}\|_2}{\|\tilde{r}\|_2} \leq \frac{\nu(\tilde{x};\theta)}{\mu(\tilde{x};\theta)} \leq \frac{\sqrt{5}+1}{2} \approx 1.618.$$

Grcar [2003] showed that ν is an asymptotically exact estimate of μ, i.e., $\nu(\tilde{x};\theta) \sim \mu(\tilde{x};\theta)$ as $\tilde{x} \to \hat{x}$.

The backward error for LS problems V

Restricting the backward perturbation to A or to b only:

- **Perturbing A:**

$$\min\{\|E\|_F;\ (A+E)^T[b-(A+E)\tilde{x}]\} = \lim_{\theta\to\infty} \mu(\tilde{x};\theta).$$

- **Perturbing b:**

$$\min\{\|f\|_2;\ A^T[(b+f)-A\tilde{x}]=0\} = \|P_A\tilde{r}\|_2 = \lim_{\theta\to 0} \frac{\mu(\tilde{x};\theta)}{\theta}.$$

See Arioli and Gratton [2008] for an application in linear regression.

Stopping criteria based on the backward error

The backward error has an important application in rounding error analysis, see, e.g., Higham [1996].

It is also recommended for construction of “general purpose” stopping criteria for iterative solvers; see, e.g., Arioli, Duff, and Ruiz [1992], Arioli, Noulard, and Russo [2001], Paige and Strakoš [2002].

Stopping criteria based on the backward error

The backward error has an important application in rounding error analysis, see, e.g., Higham [1996].

It is also recommended for construction of "general purpose" stopping criteria for iterative solvers; see, e.g., Arioli, Duff, and Ruiz [1992], Arioli, Noulard, and Russo [2001], Paige and Strakoš [2002].

Let $\tilde{x}$ be an approximation to the solution $\hat{x}$ computed by an iterative method. We say that $\tilde{x}$ is an *acceptable solution* if it solves exactly a nearby problem within a specified range of relative errors in the data. In particular, if for given α and β there exists backward perturbations E and f such that

$$(A+E)^T[(b+f)-(A+E)\tilde{x}]=0, \qquad \|E\|_F \leq \alpha\|A\|_F, \quad \|f\|_2 \leq \beta\|b\|_2.$$

Stopping criteria based on the backward error

The backward error has an important application in rounding error analysis, see, e.g., Higham [1996].

It is also recommended for construction of "general purpose" stopping criteria for iterative solvers; see, e.g., Arioli, Duff, and Ruiz [1992], Arioli, Noulard, and Russo [2001], Paige and Strakoš [2002].

Let $\tilde{x}$ be an approximation to the solution $\hat{x}$ computed by an iterative method. We say that $\tilde{x}$ is an *acceptable solution* if it solves exactly a nearby problem within a specified range of relative errors in the data. In particular, if for given α and β there exists backward perturbations E and f such that

$$(A+E)^T[(b+f)-(A+E)\tilde{x}] = 0, \qquad \|E\|_F \le \alpha\|A\|_F, \quad \|f\|_2 \le \beta\|b\|_2.$$

This is equivalent to the condition $\xi(\tilde{x};\alpha,\beta) \le 1$ which is satisfied if (and "almost only if" up to $\sqrt{2}$)

$$\mu(\tilde{x};\theta) \le \alpha\|A\|_F, \qquad \theta = \frac{\alpha\|A\|_F}{\beta\|b\|_2}.$$

See also Chang, Paige, and Titley-Peloquin [2009].

Outline

The LS residual, normal equations, and CG

Let $\tilde{x}$ be arbitrary, $\hat{x}$ be the solution of (LS) and $\tilde{r} = b - A\tilde{x}$ and $\hat{r} = b - A\hat{x}$ be the corresponding residuals.

Then the norm of the residual $\tilde{r}$ can be decomposed as

$$\|\tilde{r}\|_2^2 = \|\hat{r}\|_2^2 + \|P_A\tilde{r}\|_2^2,$$
$$\|P_A\tilde{r}\|_2 = \|P_A(b - A\tilde{x})\|_2 = \|A(\hat{x} - \tilde{x})\|_2 = \|\hat{x} - \tilde{x}\|_{A^TA}.$$

The LS residual, normal equations, and CG

Let $\tilde{x}$ be arbitrary, $\hat{x}$ be the solution of (LS) and $\tilde{r} = b - A\tilde{x}$ and $\hat{r} = b - A\hat{x}$ be the corresponding residuals.

Then the norm of the residual $\tilde{r}$ can be decomposed as

$$\|\tilde{r}\|_2^2 = \|\hat{r}\|_2^2 + \|P_A\tilde{r}\|_2^2,$$
$$\|P_A\tilde{r}\|_2 = \|P_A(b - A\tilde{x})\|_2 = \|A(\hat{x} - \tilde{x})\|_2 = \|\hat{x} - \tilde{x}\|_{A^TA}.$$

A natural way how to solve the LS problem iteratively is hence to minimize the last term which only depends on $\tilde{x}$. This is exactly what CG of Hestenes and Stiefel [1952] applied on the system of normal equations $A^TAx = A^Tb$ does!

One may apply CG directly to $A^TAx = A^Tb$ or more cleverly in the implicit way as it is done in CGLS (see Paige and Saunders [1982a]).

However, one may also use the structure of the system matrix A^TA and use another fundamental algorithm: the iterative bidiagonalization by Golub and Kahan [1965] as it is done in the LSQR algorithm by Paige and Saunders [1982a,b].

The Golub-Kahan iterative bidiagonalization

Starting with $\beta_1 u_1 = b$ and $\alpha_1 v_1 = A^T u_1$, the Golub-Kahan bidiagonalizations generates the sequence of positive α_i and β_i and orthonormal vectors u_i and v_i by two recurrent formulas

$$\beta_{k+1} u_{k+1} = A v_k - \alpha_k u_k, \qquad \alpha_{k+1} v_{k+1} = A^T u_{k+1} - \beta_{k+1} v_k,$$

for $k = 1, 2, \ldots$, where α_i and β_i are computed such that $\|u_i\|_2 = \|v_i\|_2 = 1$.

The Golub-Kahan iterative bidiagonalization

Starting with $\beta_1 u_1 = b$ and $\alpha_1 v_1 = A^T u_1$, the Golub-Kahan bidiagonalizations generates the sequence of positive α_i and β_i and orthonormal vectors u_i and v_i by two recurrent formulas

$$\beta_{k+1} u_{k+1} = A v_k - \alpha_k u_k, \qquad \alpha_{k+1} v_{k+1} = A^T u_{k+1} - \beta_{k+1} v_k,$$

for $k = 1, 2, \ldots$, where α_i and β_i are computed such that $\|u_i\|_2 = \|v_i\|_2 = 1$.

Let $U_{k+1} = [u_1, \ldots, u_k]$, $V_{k+1} = [v_1, \ldots, v_{k+1}]$, and

$$B_k = \begin{bmatrix} \alpha_1 & & \\ \beta_2 & \ddots & \\ & \ddots & \alpha_k \\ & & \beta_{k+1} \end{bmatrix}, \qquad \overline{B}_k = [B_k, \alpha_{k+1} e_{k+1}].$$

Then we have

$$U_{k+1}(\beta_1 e_1) = b, \qquad A V_k = U_{k+1} B_k, \qquad A^T U_{k+1} = V_{k+1} \overline{B}_k^T.$$

In addition, $\mathcal{R}(V_k) = \mathcal{K}_k(A^T A, A^T b)$ and $\mathcal{R}(U_k) = \mathcal{K}_k(A A^T, b)$.

The reduced LS problem in LSQR I

Let $x \in \mathcal{K}_k(A^T A, A^T b)$, $x = V_k y$, and $r = b - Ax = U_{k+1}(\beta_1 e_1 - B_k y)$. The LSQR approximation x_k is constructed by minimizing the 2-norm of the residual, i.e.,

$$\|\beta_1 e_1 - B_k y_k\|_2 = \min_y \|\beta_1 e_1 - B_k y\|_2. \qquad \text{(LSQR-LS)}$$

Hence the LSQR is equivalent to CG applied to $A^T A x = A^T b$.

The reduced LS problem in LSQR I

Let $x \in \mathcal{K}_k(A^T A, A^T b)$, $x = V_k y$, and $r = b - Ax = U_{k+1}(\beta_1 e_1 - B_k y)$. The LSQR approximation x_k is constructed by minimizing the 2-norm of the residual, i.e.,

$$\|\beta_1 e_1 - B_k y_k\|_2 = \min_y \|\beta_1 e_1 - B_k y\|_2. \qquad \text{(LSQR-LS)}$$

Hence the LSQR is equivalent to CG applied to $A^T A x = A^T b$.

The LS problem (LSQR-LS) can be solved by performing the QR decomposition of B_k

$$Q_k[B_k, \beta_1 e_1] = \begin{bmatrix} R_k & f_k \\ 0 & \bar{\phi}_{k+1} \end{bmatrix}.$$

The action of the kth reflection can be described as

$$\begin{bmatrix} c_k & s_k \\ s_k & -c_k \end{bmatrix} \begin{bmatrix} \bar{\rho}_k & 0 & \bar{\phi}_k \\ \beta_{k+1} & \alpha_{k+1} & 0 \end{bmatrix} = \begin{bmatrix} \rho_k & \theta_{k+1} & \phi_k \\ 0 & \bar{\rho}_{k+1} & \bar{\phi}_{k+1} \end{bmatrix}.$$

The vector of coordinates is then obtained by solving for y_k the system

$$R_k y_k = f_k \equiv [\phi_1, \ldots, \phi_k]^T.$$

The reduced LS problem in LSQR II

The residuals $r_k \equiv b - Ax_k$ and $t_k \equiv \beta_1 e_1 - B_k y_k$ then satisfy

$$r_k = U_{k+1} t_k, \qquad t_k = \bar{\phi}_{k+1} Q_k^T e_{k+1}, \qquad \|r_k\|_2 = \bar{\phi}_{k+1}.$$

In addition,

$$\bar{\phi}_k^2 = \phi_{k+1}^2 + \bar{\phi}_{k+1}^2 \qquad \Rightarrow \qquad \|r_{k-1}\|_2^2 - \|r_k\|_2^2 = \phi_k^2.$$

Note that a recurrent formula for the approximations can be derived by computing direction vectors $W_k = [w_1, \ldots, w_k]$ such that

$$x_k = V_k y_k = V_k R_k^{-1} f_k = W_k f_k.$$

Stopping criteria in LSQR

In the original implementation of LSQR, see Paige and Saunders [1982a], the stopping criteria are based on ξ_{CS} and on E_1 of Stewart [1977]. The iteration is terminated when

$$\|r_k\|_2 \leq \alpha\|A\|_F\|x_k\|_2 + \beta\|b\|_2$$

or

$$\frac{\|A^T r_k\|_2}{\|r_k\|_2} \leq \alpha\|A\|_F.$$

Both $\|r_k\|_2$ and $\|A^T r_k\|_2$ can be cheaply computed in LSQR.

It was observed by Chang, Paige, and Titley-Peloquin [2009] that the stopping criterion using $A^T r_k$ can be very conservative; moreover, the convergence of $\|A^T r_k\|_2$ to zero can be very irregular (see also Fong and Saunders [2010] for LSMR).

Stopping criteria in LSQR: Numerical illustration I

Least squares problem ASH608 from the Harwell Boeing collection

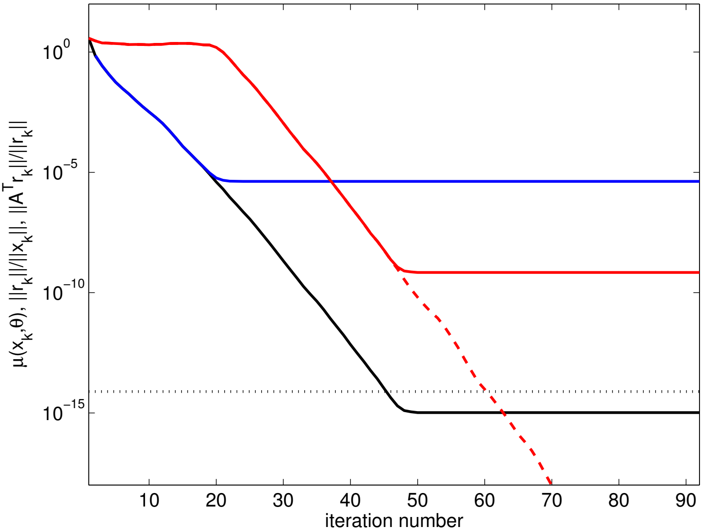

Stopping criteria in LSQR: Numerical illustration II

Least squares problem ILLC1033 from the Harwell Boeing collection

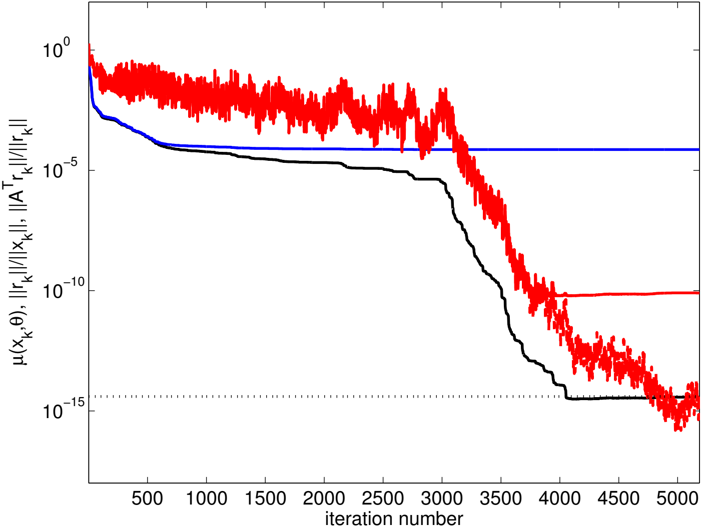

Outline

Our goal is to provide estimates of

- the norm of the projected residual $P_A r_k$,
- the backward error $\mu(x_k;\theta)$,
- the asymptotic estimate $\nu(x_k;\theta)$ of $\mu(x_k;\theta)$,

which are cheaply computable in LSQR.

Our goal is to provide estimates of

- the norm of the projected residual $P_A r_k$,
- the backward error $\mu(x_k;\theta)$,
- the asymptotic estimate $\nu(x_k;\theta)$ of $\mu(x_k;\theta)$,

which are cheaply computable in LSQR.

The norm of $P_A\tilde{r}$ can be used to provide a bound

$$\mu(\tilde{x};\theta) \le \frac{\theta\|P_A\tilde{r}\|_2}{\sqrt{1+\theta^2\|\tilde{x}\|_2^2}},$$

which is similar to the backward perturbation $E_2 = (\hat{r}-\tilde{r})\tilde{x}^\dagger$ provided by Stewart [1977]; see Chang, Paige, and Titley-Peloquin [2009].

Our goal is to provide estimates of

- the norm of the projected residual $P_A r_k$,
- the backward error $\mu(x_k; \theta)$,
- the asymptotic estimate $\nu(x_k; \theta)$ of $\mu(x_k; \theta)$,

which are cheaply computable in LSQR.

The norm of $P_A \tilde{r}$ can be used to provide a bound

$$\mu(\tilde{x}; \theta) \leq \frac{\theta \|P_A \tilde{r}\|_2}{\sqrt{1 + \theta^2 \|\tilde{x}\|_2^2}},$$

which is similar to the backward perturbation $E_2 = (\hat{r} - \tilde{r})\tilde{x}^\dagger$ provided by Stewart [1977]; see Chang, Paige, and Titley-Peloquin [2009].

Malyshev and Sadkane [2002] estimate the backward error by running the iterative Golub-Kahan bidiagonalization on A and the residual vector $\tilde{r}$. In Grcar, Saunders, and Su [2007] the backward error estimates are based on ν; for sparse LS problems solved by LSQR an inner iteration is applied to estimate the backward error.

Estimating $\|P_A r_k\|_2$ I

LSQR is equivalent to CG applied to $A^T Ax = A^T b$, so it minimizes $\|\hat{x} - x_k\|_{A^T A}$ at each step. We have

$$\|\hat{x} - x_k\|_{A^T A} = \|A(\hat{x} - x_k)\|_2 = \|P_A(b - Ax_k)\|_2 = \|P_A r_k\|_2,$$

so $\|P_A r_k\|_2$ has the meaning of the energy norm in the equivalent CG on $A^T Ax = A^T b$.

The extensive amount of work was devoted to the estimation of the error in CG, see, e.g., Meurant and Strakoš [2006] and the references therein. Here we concentrate on the lower bound proposed by Hestenes and Stiefel [1952]. It use in numerical computations was justified by Strakoš and Tichý [2002].

Estimating $\|P_A r_k\|_2$ II

Let d be a positive integer. We have

$$\|r_k\|_2^2 = \|P_A r_k\|_2^2 + \|\hat{r}\|_2^2, \qquad \|r_{k+d}\|_2^2 = \|P_A r_{k+d}\|_2^2 + \|\hat{r}\|_2^2.$$

This gives

$$\begin{aligned} \|P_A r_k\|_2^2 &= \|r_k\|_2^2 - \|r_{k+d}\|_2^2 + \|P_A r_{k+d}\|_2^2 \\ &= \sum_{i=k+1}^{k+d} (\|r_{i-1}\|_2^2 - \|r_i\|_2^2) + \|P_A r_{k+d}\|_2^2 \\ &= \sum_{i=k+1}^{k+d} \phi_i^2 + \|P_A r_{k+d}\|_2^2 \end{aligned}$$

Estimating $\|P_A r_k\|_2$ III

Let

$$\lambda_d(x_k) \equiv \sum_{i=k+1}^{k+d} \phi_i^2.$$

We can use $\lambda_d(x_k)$ to estimate

- the norm of $P_A r_k$, i.e., the minimal backward perturbation of the right-hand side in the least squares problem (LS),
- the backward error $\mu(x_k; \theta)$ by

$$\mu(x_k; \theta) \approx \frac{\theta \lambda_d(x_k)}{\sqrt{1 + \theta^2 \|x_k\|_2^2}}.$$

Estimating $\|P_A r_k\|_2$ III

Let

$$\lambda_d(x_k) \equiv \sum_{i=k+1}^{k+d} \phi_i^2.$$

We can use $\lambda_d(x_k)$ to estimate

- the norm of $P_A r_k$, i.e., the minimal backward perturbation of the right-hand side in the least squares problem (LS),
- the backward error $\mu(x_k;\theta)$ by

$$\mu(x_k;\theta) \approx \frac{\theta\lambda_d(x_k)}{\sqrt{1+\theta^2\|x_k\|_2^2}}.$$

The estimate $\lambda_d(x_k)$ satisfies

$$\lambda_d(x_k) = \|B_{k+d}B_{k+d}^{\dagger}\bar{t}_k\|_2, \qquad \bar{t}_k = \beta_1 e_1 - B_{k+d}\begin{bmatrix} y_k \\ 0 \end{bmatrix},$$

so it can be interpreted as the norm of the minimal backward perturbation of the right-hand side in the LS problem solved by LSQR at step $k+d$.

Estimating $\|P_A r_k\|_2$: Numerical illustration

Least squares problem WELL1033 from the Harwell Boeing collection

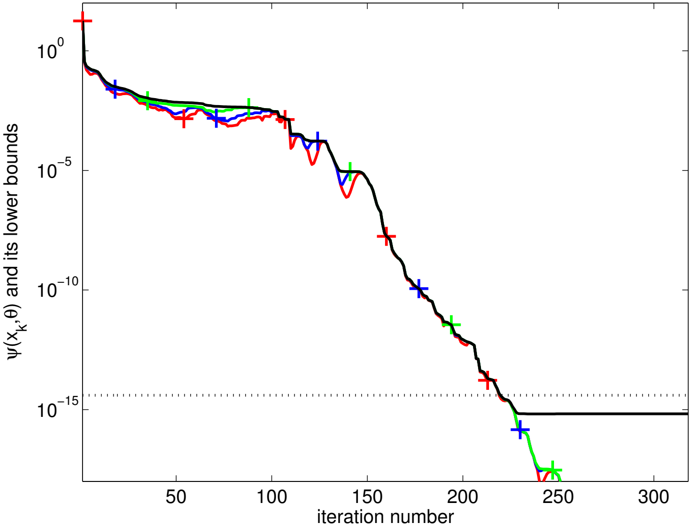

Estimating $\mu(x_k;\theta)$ I

We wish to estimate

$$\mu(x_k;\theta) = \min\{\|[E,\theta f]\|_F;\ (A+E)^T[(b+f)-(A+E)x_k]=0\} \\ \equiv g(x_k,\theta,A,b).$$

Let

$$\underline{\mu}_d(x_k;\theta) \equiv g\left(\begin{bmatrix} y_k \\ 0 \end{bmatrix}, \theta, B_{k+d}, \beta_1 e_1\right), \quad \overline{\mu}_d(x_k;\theta) \equiv g\left(\begin{bmatrix} y_k \\ 0 \end{bmatrix}, \theta, \overline{B}_{k+d}, \beta_1 e_1\right).$$

The estimate $\underline{\mu}_d(x_k;\theta)$ can be (similarly as $\lambda_d(x_k)$) interpreted as the backward error of the kth approximation of LSQR in the reduced least squares problem solved at step $k+d$.

The estimates $\underline{\mu}_d(x_k;\theta)$ and $\overline{\mu}_d(x_k;\theta)$ represent lower and upper bounds for $\mu(x_k;\theta)$ which become tighter with increasing d.

Estimating $\mu(x_k;\theta)$ II

In order to evaluate the estimates we have to compute the minimal singular values of the matrices

$$N_{k,d} \equiv \left[B_{k+d}, \omega_k(I - \bar{t}_k \bar{t}_k^{\dagger})\right], \qquad \overline{N}_{k,d} \equiv \left[\overline{B}_{k+d}, \omega_k(I - \bar{t}_k \bar{t}_k^{\dagger})\right],$$

where $\omega_k \equiv \theta\|r_k\|_2/(1+\theta^2\|x_k\|_2^2)^{1/2}$, $\bar{t}_k = [t_k^T, 0^T]^T$.

Estimating $\mu(x_k;\theta)$ II

In order to evaluate the estimates we have to compute the minimal singular values of the matrices

$$N_{k,d} \equiv \left[B_{k+d}, \omega_k(I - \bar{t}_k\bar{t}_k^\dagger)\right], \qquad \overline{N}_{k,d} \equiv \left[\overline{B}_{k+d}, \omega_k(I - \bar{t}_k\bar{t}_k^\dagger)\right],$$

where $\omega_k \equiv \theta\|r_k\|_2/(1+\theta^2\|x_k\|_2^2)^{1/2}$, $\bar{t}_k = [t_k^T, 0^T]^T$.

However, the vector $t_k = U_k^T r_k$ is not available in LSQR. We can use the relation $t_k = \bar{\phi}_{k+1}Q_k^T e_{k+1}$ to show that $N_{k,d}$ has the same singular values as

$$M_{k,d} \equiv [\hat{Q}_k B_{k+d}, \omega_k(I - e_{k+1}e_{k+1}^T)].$$

The matrix $M_{k,d}$ can be transformed to a bidiagonal form using $2(k+d+1)$ Givens rotations and the smallest singular value can be computed using the method of bisection.

We can proceed similarly with the matrix $\overline{N}_{k,d}$.

Estimating $\mu(x_k; \theta)$: Numerical illustration I

Least squares problem WELL1033 from the Harwell Boeing collection

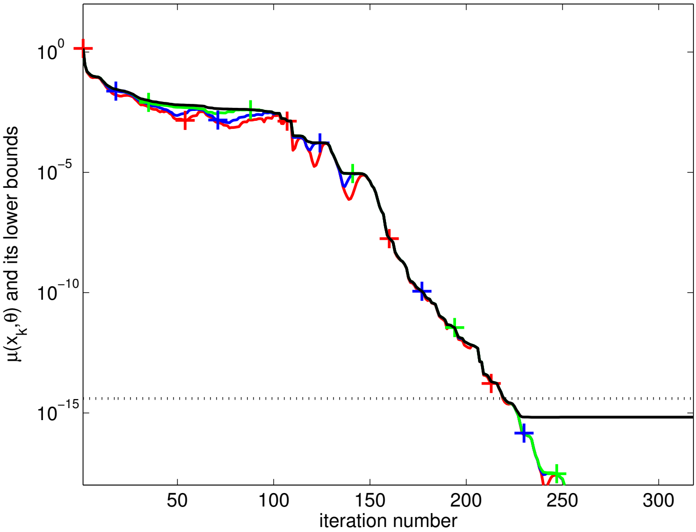

Estimating $\mu(x_k;\theta)$: Numerical illustration II

Least squares problem WELL1033 from the Harwell Boeing collection

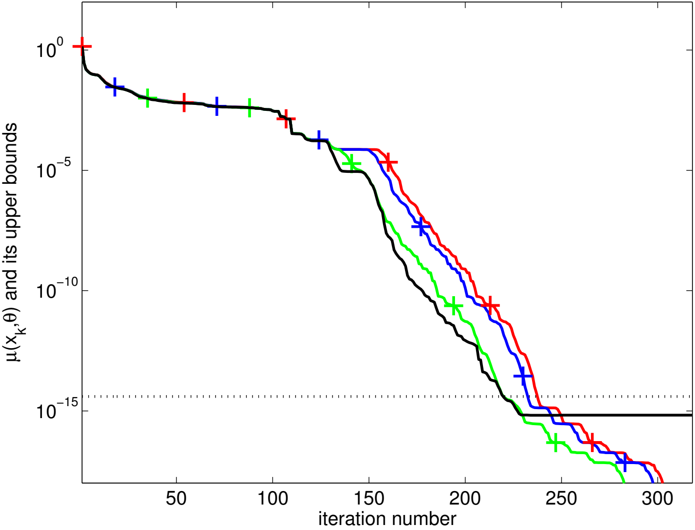

Estimating $\nu(x_k;\theta)$

Similarly to $\mu(x_k;\theta)$ we can bound the quantity $\nu(x_k;\theta)$,

$$\nu(x_k;\theta) = \frac{\omega_k}{\|r_k\|_2} \left\| \begin{bmatrix} A \\ \omega_k I \end{bmatrix} \begin{bmatrix} A \\ \omega_k I \end{bmatrix}^{\dagger} \begin{bmatrix} r_k \\ 0 \end{bmatrix} \right\|_2, \qquad \omega_k = \frac{\theta \|r_k\|_2}{\sqrt{1+\theta^2\|x_k\|_2^2}},$$

using

$$\underline{\nu}_d(x_k;\theta) = \left\| \begin{bmatrix} B_{k+d} \\ \omega_k I \end{bmatrix} \begin{bmatrix} B_{k+d} \\ \omega_k I \end{bmatrix}^{\dagger} \begin{bmatrix} \bar{t}_k \\ 0 \end{bmatrix} \right\|_2,$$

$$\overline{\nu}_d(x_k;\theta) = \left\| \begin{bmatrix} \overline{B}_{k+d} \\ \omega_k I \end{bmatrix} \begin{bmatrix} \overline{B}_{k+d} \\ \omega_k I \end{bmatrix}^{\dagger} \begin{bmatrix} \bar{t}_k \\ 0 \end{bmatrix} \right\|_2.$$

Outline

Conclusions and future work

- We presented cheaply computable estimates of the backward error in LS problems in the context of the LSQR algorithm. To some extent they can be applied also to another equivalent algorithms like CGLS.
- However, their applicability is limited, e.g., when using a preconditioner.

Conclusions and future work

- We presented cheaply computable estimates of the backward error in LS problems in the context of the LSQR algorithm. To some extent they can be applied also to another equivalent algorithms like CGLS.
- However, their applicability is limited, e.g., when using a preconditioner.

A possible outlook:

- A refined analysis of the relation between $\|P_A\tilde{r}\|_2$ and μ, especially the lower bound if possible.
- Quantitative analysis of $\underline{\mu} \approx \mu$?
- Estimation of minimal backward perturbation in related problems like damped least squares problems

$$\min_x \left\{ \|b - Ax\|_2^2 + \lambda^2 \|x\|_2^2 \right\}.$$

Thank you for your attention!

For more information, see

P.J., David Titley-Peloquin, Estimating the backward error in LSQR, *SIAM J. Matrix Anal. Appl.*, 31(4):2055–2074, 2010.

References I

M. Arioli and S. Gratton. Least-squares problems, normal equations, and stopping criteria for the conjugate gradient method. Technical Report RAL-TR-2008-008, Rutherford Appleton Laboratory, 2008.

M. Arioli, I. S. Duff, and D. Ruiz. Stopping criteria for iterative solvers. *SIAM J. Matrix Anal. Appl.*, 13(1):138–144, 1992.

M. Arioli, E. Noulard, and A. Russo. Stopping criteria for iterative methods: applications to PDE's. *Calcolo*, 38:97–112, 2001.

Å. Björck. *Numerical Methods for Least Squares Problems*. SIAM, Philadelphia, 1996.

X.-W. Chang, C. C. Paige, and D. Titley-Peloquin. Stopping criteria for the iterative solution of linear least squares problems. *SIAM J. Matrix Anal. Appl.*, 31(2):831–852, 2009.

D. C.-L. Fong and M. A. Saunders. LSMR: an iterative algorithm for sparse least-squares problems. Technical Report SOL 2010-2, Department of Management Science and Engineering, Stanford University, 2010.

References II

G. H. Golub and W. M. Kahan. Calculating the singular values and pseudo-inverse of a matrix. *SIAM J. Numer. Anal. Ser. B*, 2(2):205–224, 1965.

G. H. Golub and C. F. Van Loan. *Matrix Computations*. The Johns Hopkins University Press, Baltimore, third edition, 1996.

J. F. Grcar. Optimal sensitivity analysis of linear least squares. Technical Report LBNL-52434, Lawrence Berkeley National Laboratory, 2003.

J. F. Grcar, M. A. Saunders, and Z. Su. Estimates of optimal backward perturbations for linear least squares problems. Technical Report SOL 2007-1, Department of Management Science and Engineering, Stanford University, 2007.

M. Gu. Backward perturbation bounds for linear least squares problems. *SIAM J. Matrix Anal. Appl.*, 20(2):363–372, 1998.

M. T. Heath. Numerical methods for large sparse linear least squares problems. *SIAM J. Sci. Stat. Comput.*, 5(3):497–513, 1984.

M. R. Hestenes and E. Stiefel. Methods of conjugate gradients for solving linear systems. *J. Res. Nat. Bur. Stand.*, 49(6):409–435, 1952.

References III

N. J. Higham. *Accuracy and Stability of Numerical Algorithms*. SIAM, Philadelphia, PA, 1996.

R. Karlson and B. Waldén. Estimation of optimal backward perturbation bounds for the linear least squares problem. *BIT*, 37(4):862–869, 1997.

C. L. Lawson and R. J. Hanson. *Solving Least Squares Problems*. Prentice-Hall, Englewood Cliffs, NJ, 1974.

A. N. Malyshev and M. Sadkane. Computation of optimal backward perturbation bounds for large sparse linear least squares problems. *BIT*, 41(4):739–747, 2002.

G. Meurant and Z. Strakoš. The Lanczos and conjugate gradient algorithms in finite precision arithmetic. *Acta Numer.*, 15:471–542, 2006.

C. C. Paige and M. A. Saunders. LSQR: an algorithm for sparse linear equations and sparse least squares. *ACM Trans. Math. Software*, 8(1):43–71, 1982a.

C. C. Paige and M. A. Saunders. ALGORITHM 583; LSQR: Sparse linear equations and least-squares problems. *ACM Trans. Math. Software*, 8(2): 195–209, 1982b.

References IV

C. C. Paige and Z. Strakoš. Residual and backward error bounds in minimum residual Krylov subspace methods. *SIAM J. Sci. Comput.*, 23(6):1898–1923, 2002.

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):543–548, 1967.

G. W. Stewart. Research, development, and LINPACK. In *Mathematical Software III*, pages 1–14. Academic Press, New York, 1977.

Z. Strakoš and P. Tichý. On error estimation in the conjugate gradient method and why it works in finite precision computations. *Electron. Trans. Numer. Anal.*, 13:56–80, 2002.

B. Waldén, R. Karlson, and J.-G. Sun. Optimal backward perturbation bounds for the linear least squares problem. *Numer. Linear Algebra Appl.*, 2(3):271–286, 1995.