Perturbation Bounds for Linear Regression Problems

Bert W. Rust
Computing and Applied Mathematics Division
Building 101, Room A-238
National Institute of Standards and Technology
Gaithersburg, MD 20899
bwr@cam.nist.gov

Abstract

This paper examines errors in the estimated solution vector $\hat{\mathbf{x}}$ to the linear regression problem

$$\hat{\mathbf{y}} = \mathbf{K}\mathbf{x}^* + \hat{\epsilon} \ , \quad \mathcal{E}(\hat{\epsilon}) = \mathbf{o} \ , \quad \mathcal{E}\left(\hat{\epsilon}\hat{\epsilon}^T\right) = \mathbf{S}^2 \ ,$$

when the dominant uncertainties are the measuring errors $\hat{\epsilon}$. Backward error analysis gives the hopelessly pessimistic bound

$$\frac{\| \hat{\mathbf{x}} - \mathbf{x}^* \|_2}{\| \mathbf{x}^* \|_2} \leq \mathrm{cond}(\mathbf{S}^{-1}\mathbf{K}) \frac{\| \mathbf{S}^{-1}\hat{\epsilon} \|_2}{\| \mathbf{S}^{-1}\mathbf{K}\mathbf{x}^* \|_2} \ ,$$

by assuming the worst possible combination of random errors, an extremely unlikely occurence for nontrivial problems. A statistical treatment yields a more realistic bound on the expected uncertainty in a single element $\hat{x}_i$ which does not depend on $\mathrm{cond}(\mathbf{S}^{-1}\mathbf{K})$. Classical regression theory provides easily computable confidence intervals for the individual $\hat{x}_i$ separately.

Notation and Test Problem

Statisticians write the $m \times n$ linear regression model as

$$\mathbf{Y} = \mathbf{X}\boldsymbol{\beta} + \epsilon \ , \ \ \mathcal{E}(\epsilon) = \mathbf{o} \ , \ \ \mathcal{E}\left(\epsilon\,\epsilon^T\right) = \mathbf{\Sigma}^2 \ , \qquad (1)$$

where $\mathbf{Y}$ is a measured m-vector containing measuring errors ϵ, $\mathbf{X}$ is a known $m \times n$ matrix with $m \geq n = \mathrm{rank}(\mathbf{X})$, and $\boldsymbol{\beta}$ is the vector to be estimated. Numerical analysts write the linear least squares problem as

$$\rho_{LS}^2 = \min_{x \in R^n} ||b - Ax||_2^2 \ , \qquad (2)$$

where b is the measured m-vector, A is the $m \times n$ matrix, x is the vector to be estimated, $||b - Ax||_2^2$ is the squared two-norm of the residual vector, and ρ_{LS}^2 is the minimum sum of squared residuals. They usually assume (but seldom state) the linear regression model

$$b = Ax^* + \delta b \ , \ \ \mathcal{E}(\delta b) = 0 \ , \ \ \mathcal{E}\left(\delta b\, \delta b^T\right) = \sigma^2 I_m \ , \qquad (3)$$

where I_m is the mth order identity matrix, and the scalar σ is unknown.

Since choosing either of the above notations would deeply offend one of the two schools, consider

$$\hat{\mathbf{y}} = \mathbf{K}\mathbf{x}^* + \hat{\epsilon} \ , \ \ \mathcal{E}(\hat{\epsilon}) = \mathbf{o} \ , \ \ \mathcal{E}\left(\hat{\epsilon}\,\hat{\epsilon}^T\right) = \mathbf{S}^2 \ , \qquad (4)$$

where $\hat{\mathbf{y}}$ is the measured m-vector, and $\mathbf{K}$ is the known $m \times n$ matrix with $\mathrm{rank}(\mathbf{K}) = n$. This notation is appropriate when linear regression is applied to systems of integral equations of the form

$$\hat{y}_i = \int_a^b K_i(\xi) x(\xi)\, d\xi + \hat{\epsilon}_i \ , \ \ i = 1, 2, \ldots, m \ , \qquad (5)$$

where the $\hat{y}_i$ are measured values, the $K_i(\xi)$ are known functions, and $x(\xi)$ is the function to be estimated. Such equations are widely used to model the effects of a measuring instrument on the thing being measured. One way to approximate $x(\xi)$ is to replace the integrals with quadrature sums, i.e.,

$$\int_a^b K_i(\xi) x(\xi) d\xi \approx \sum_{j=1}^{n} \omega_j K_i(\xi_j) x(\xi_j) \ , \qquad (6)$$

where the ω_j are prescribed quadrature coefficients and the $x(\xi_j)$ form a discrete approximation to $x(\xi)$. It is important to choose n large enough so that the quadrature errors are small relative to the $\hat{\epsilon}_i$. If the sums are substituted for the integrals in (5) and the products $\omega_j K_i(\xi_j)$ collected into a matrix $\mathbf{K}$, the result is the model (4).

A test problem capturing many of the salient features of real instrument correction problems is obtained by discretizing the Phillips [5] equation

$$y(t) = \int_{-3}^{3} K(t, \xi) x(\xi)\, d\xi \ , \ \ -6 \leq t \leq 6 \ , \qquad (7)$$

with

$$K(t, \xi) = \begin{cases} 1 + \cos\left[\frac{\pi(\xi - t)}{3}\right] \ , & | \xi - t | \leq 3 \\ & | t | \leq 6 \\ 0 \ , \ \text{otherwise} \ , & \end{cases} \qquad (8)$$

and

$$y(t) = \begin{cases} (6-\mid t \mid)\left[1+\frac{1}{2}\cos\left(\frac{\pi t}{3}\right)\right] & \\ \quad + \frac{9}{2\pi}\sin\left(\frac{\pi|t|}{3}\right), & \mid t \mid\leq 6 \\ 0\ ,\ \text{otherwise}\ . & \end{cases} \tag{9}$$

The kernel $K(t,\xi)$ is non-negative, with maximum value 2, attained on the line $t=\xi$. The solution is

$$x(\xi) = \begin{cases} 1+\cos\left(\frac{\pi\xi}{3}\right) & , \quad \mid \xi \mid\leq 3 \\ 0 & , \quad \text{otherwise} \end{cases} \tag{10}$$

The functions $y(t)$ and $x(\xi)$ are plotted in Figure 1.

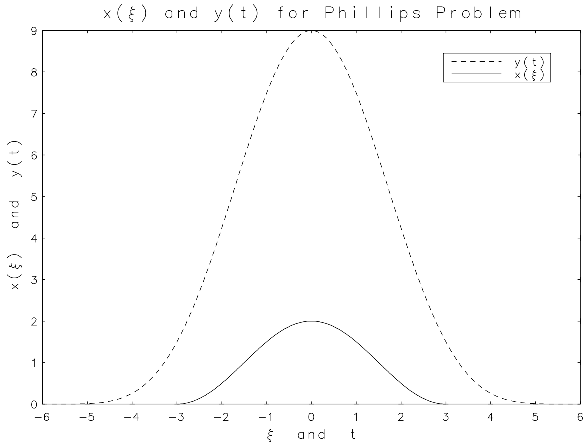

Figure 1:

Discretizing replaces continuous variables t and ξ with meshes $t_i, i=1,\ldots,m$ and $\xi_j, j=1,\ldots,n$. Choosing $m=150$ equi-spaced t_i on $-5.925\leq t\leq 5.925$ and using an $n=121$ point trapezoidal rule on $-3.0\leq\xi\leq 3.0$ gave

$$\mathbf{y}^* \equiv \mathbf{K}\mathbf{x}^*\ , \tag{11}$$

where $\mathbf{x}^*$ is a 121-vector of $x(\xi_j)$ computed by (10), and $\mathbf{y}^*$ was computed by (11) rather than (9) to assure that the $\hat{\epsilon_i}$ were the only errors in the model. The $\hat{\epsilon_i}$ were obtained by random sampling from $N(\mathbf{o},\mathbf{S}^2)$ with

$$\mathbf{S} = \mathbf{diag}(s_1,\ s_2,\ \ldots,\ s_m)\ ,\ \ s_i=(10^{-6})y_i^*\ , \tag{12}$$

which means that the errors in the $\hat{y}_i$ were in the 6th digit. The discretized model can thus be written

$$\mathbf{y}^* \equiv \mathbf{K}\mathbf{x}^*\ ,\ \ \hat{\mathbf{y}} = \mathbf{K}\mathbf{x}^* + \hat{\boldsymbol{\epsilon}}\ ,\ \ \hat{\boldsymbol{\epsilon}} \sim N(\ \mathbf{o}\ ,\ \mathbf{S}^2\ )\ , \tag{13}$$

and the least squares estimate

$$\hat{\mathbf{x}} = \left(\mathbf{K}^T\mathbf{S}^{-2}\mathbf{K}\right)^{-1}\mathbf{K}^T\mathbf{S}^{-2}\hat{\mathbf{y}}\ , \tag{14}$$

computed by LINPACK subroutines DQRDC and DQRSL [2], is shown in Figure 2. The dashed curve

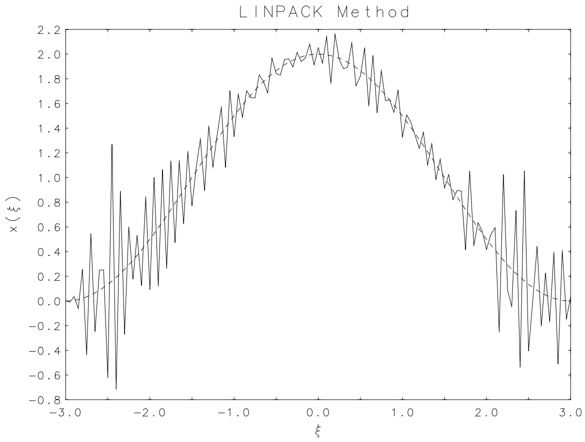

Figure 2:

is $x(t)$ and the jagged curve is the estimate. The large oscillations are induced by errors in the 6th digit of the $\hat{y}_i$! Such ill-conditioning is typical of regression models arising from discretized first kind integral equations.

Classical Perturbation Theory

To simplify the discussion in this section, let

$$\hat{\mathbf{b}} \equiv \mathbf{S}^{-1}\hat{\mathbf{y}}\ ,\ \ \mathbf{A}\equiv\mathbf{S}^{-1}\mathbf{K}\ ,\ \ \delta\mathbf{b}\equiv\mathbf{S}^{-1}\hat{\boldsymbol{\epsilon}}\ , \tag{15}$$

and rewrite (13) as

$$\mathbf{b}^* = \mathbf{A}\mathbf{x}^*\ ,\ \ \hat{\mathbf{b}} = \mathbf{A}\mathbf{x}^* + \delta\mathbf{b}\ ,\ \ \delta\mathbf{b}\sim N(\ \mathbf{o}\ ,\ \mathbf{I}_m\ )\ . \tag{16}$$

The problem of interest is to find bounds for the errors in the least squares solution $\hat{\mathbf{x}} = (\mathbf{A}^T\mathbf{A})^{-1}\mathbf{A}^T\hat{\mathbf{b}}$.

The traditional approach ignores $\mathbf{x}^*$ and the statistical assumptions about $\delta\mathbf{b}$, seeking instead to bound the difference between estimates corresponding to two different $\hat{\mathbf{b}}$ vectors. One of these, $\mathbf{b}$, corresponds to the problem

$$||\mathbf{A}\mathbf{x}-\mathbf{b}||_2 = \min = \rho_{LS}\ , \tag{17}$$

and the other, $\mathbf{b}+\Delta\mathbf{b}$, corresponds to a perturbed problem

$$||(\mathbf{A}+\Delta\mathbf{A})\hat{\mathbf{x}} - (\mathbf{b}+\Delta\mathbf{b})||_2 = \min\ , \tag{18}$$

where $\Delta\mathbf{b}$ and $\Delta\mathbf{A}$ represent the uncertainties in $\mathbf{b}$ and $\mathbf{A}$. The regression model assumes that $\mathbf{A}$ is known exactly, or at least to much higher precision than $\mathbf{b}$, but numerical analysts argue that truncation errors arising when $\mathbf{A}$ is read into a finite-accuracy computer should be taken into account. A long and intricate argument [3] leads to the following error bound:

$$\frac{||\hat{\mathbf{x}}-\mathbf{x}||_2}{||\mathbf{x}||_2} \leq \varepsilon\left\{\frac{2\kappa(\mathbf{A})||\mathbf{b}||_2 + \rho_{LS}[\kappa(\mathbf{A})]^2}{\sqrt{||\mathbf{b}||_2^2-\rho_{LS}^2}}\right\} + \mathcal{O}(\varepsilon^2)\ , \tag{19}$$

where

$$\varepsilon = \max\left\{\frac{\|\Delta\mathbf{A}\|_2}{\|\mathbf{A}\|_2}\ ,\ \frac{\|\Delta\mathbf{b}\|_2}{\|\mathbf{b}\|_2}\right\}\ , \tag{20}$$

and

$$\kappa(\mathbf{A}) = \mathrm{cond}(\mathbf{A}) = \frac{\sigma_{\max}(\mathbf{A})}{\sigma_{\min}(\mathbf{A})} = \frac{\sigma_1}{\sigma_n} \tag{21}$$

is the *condition number* which is just the ratio of the largest to the smallest singular value of $\mathbf{A}$.

While numerical analysts are fascinated by the truncation $\Delta\mathbf{A}$, people who actually make measurements usually insist on a computer arithmetic with enough precision to render such perturbations negligible in comparision to the measurement errors. When the Computer Acquisition Committee at the National Bureau of Standards was writing specifications for a new computer in 1984, some members insisted on a machine with 64-bit single precision because 32-bit machines give only 6 to 7 digits of precision, and they routinely measured things better than that. Accordingly, let $\Delta\mathbf{A} = \mathbf{0}$. This leads to the more easily obtained [6] bound

$$\frac{\|\hat{\mathbf{x}} - \mathbf{x}\|_2}{\|\mathbf{x}\|_2} \le \mathrm{cond}(\mathbf{A})\frac{\|\Delta\mathbf{b}\|_2}{\|\mathbf{b}\|_2}\ , \tag{22}$$

which also depends strongly on $\mathrm{cond}(\mathbf{A})$.

Assessing the Classical Bound

The bound (22) is computable, but it does not relate a computed estimate to $\mathbf{x}^*$. To obtain such a result, let

$$\mathbf{b} = \mathbf{b}^* = \mathbf{A}\mathbf{x}^*\ ,\ \ \Delta\mathbf{b} = \delta\mathbf{b} \sim N(\mathbf{o}, \mathbf{I}_m)\ , \tag{23}$$

and replace problems (17) and (18) with

$$\|\mathbf{A}\mathbf{x}^* - \mathbf{b}^*\|_2 = \min = 0\ ,\ \ \|\mathbf{A}\hat{\mathbf{x}} - (\mathbf{b}^* + \delta\mathbf{b})\|_2 = \min\ . \tag{24}$$

The bound (22) then becomes

$$\frac{\|\hat{\mathbf{x}} - \mathbf{x}^*\|_2}{\|\mathbf{x}^*\|_2} \le \mathrm{cond}(\mathbf{A})\frac{\|\delta\mathbf{b}\|_2}{\|\mathbf{A}\mathbf{x}^*\|_2}\ , \tag{25}$$

which is not practicable because it depends on $\mathbf{x}^*$. But $\mathbf{x}^*$ is known for the test problem, and this provides a means for evaluating the perturbation bound. To restore the original notation, substitute (15) into (25) to obtain

$$\frac{\|\hat{\mathbf{x}} - \mathbf{x}^*\|_2}{\|\mathbf{x}^*\|_2} \le \mathrm{cond}(\mathbf{S}^{-1}\mathbf{K})\frac{\|\mathbf{S}^{-1}\hat{\epsilon}\|_2}{\|\mathbf{S}^{-1}\mathbf{K}\mathbf{x}^*\|_2}\ , \tag{26}$$

where

$$\mathrm{cond}(\mathbf{S}^{-1}\mathbf{K}) = \frac{\sigma_{\max}(\mathbf{S}^{-1}\mathbf{K})}{\sigma_{\min}(\mathbf{S}^{-1}\mathbf{K})} = \frac{\sigma_1}{\sigma_n}\ . \tag{27}$$

Multiplying (26) by $\|\mathbf{x}^*\|_2$ and squaring both sides gives

$$\|\hat{\mathbf{x}} - \mathbf{x}^*\|_2^2 \le \frac{\left[\mathrm{cond}(\mathbf{S}^{-1}\mathbf{K})\right]^2 \|\mathbf{x}^*\|_2^2}{\|\mathbf{S}^{-1}\mathbf{K}\mathbf{x}^*\|_2^2}\,\|\mathbf{S}^{-1}\hat{\epsilon}\|_2^2\ . \tag{28}$$

Since both sides are non-negative functions of the random vector $\hat{\epsilon}$, it follows that

$$\mathcal{E}\left(\|\hat{\mathbf{x}} - \mathbf{x}^*\|_2^2\right) \le \frac{\left[\mathrm{cond}(\mathbf{S}^{-1}\mathbf{K})\right]^2 \|\mathbf{x}^*\|_2^2}{\|\mathbf{S}^{-1}\mathbf{K}\mathbf{x}^*\|_2^2}\,\mathcal{E}\left(\|\mathbf{S}^{-1}\hat{\epsilon}\|_2^2\right)\ . \tag{29}$$

It follows from (13) that $\mathbf{S}^{-1}\hat{\epsilon} \sim N(\ \mathbf{o}\ ,\ \mathbf{I}_m\ )$ which implies $\|\mathbf{S}^{-1}\hat{\epsilon}\|_2^2 \sim \chi^2(m)$, so $\mathcal{E}\left(\|\mathbf{S}^{-1}\hat{\epsilon}\|_2^2\right) = m$. Therefore

$$\mathcal{E}\left(\|\hat{\mathbf{x}} - \mathbf{x}^*\|_2^2\right) \le \frac{m\left[\mathrm{cond}(\mathbf{S}^{-1}\mathbf{K})\right]^2 \|\mathbf{x}^*\|_2^2}{\|\mathbf{S}^{-1}\mathbf{K}\mathbf{x}^*\|_2^2}\ , \tag{30}$$

which relates $\hat{\mathbf{x}}$ to $\mathbf{x}^*$, but with the elements of $|\hat{\mathbf{x}} - \mathbf{x}^*|$ muddled together. To clarify, define $|\Delta x|_{rms}$ by

$$|\Delta x|_{rms}^2 \equiv \mathcal{E}\left(\frac{1}{n}\sum_{j=1}^{n}|\hat{x}_j - x_j^*|^2\right) = \frac{1}{n}\mathcal{E}\left(\|\hat{\mathbf{x}} - \mathbf{x}^*\|_2^2\right)\ , \tag{31}$$

so by (30),

$$|\Delta x|_{rms} \le \left(\sqrt{\frac{m}{n}}\right)\mathrm{cond}(\mathbf{S}^{-1}\mathbf{K})\frac{\|\mathbf{x}^*\|_2}{\|\mathbf{S}^{-1}\mathbf{K}\mathbf{x}^*\|_2}\ . \tag{32}$$

The quantity $|\Delta x|_{rms}$ is the expected root mean squared absolute error for the components of $\hat{\mathbf{x}}$. The test problem has $\|\mathbf{x}^*\|_2 = 13.82$, $\sigma_1(\mathbf{S}^{-1}\mathbf{K}) = 3.3950 \times 10^9$, and $\sigma_{121}(\mathbf{S}^{-1}\mathbf{K}) = 1.1610$. Thus $\mathrm{cond}(\mathbf{S}^{-1}\mathbf{K}) = 2.924 \times 10^9$, and by (12),

$$\mathbf{S}^{-1}\mathbf{K}\mathbf{x}^* = \mathbf{S}^{-1}\mathbf{y}^* = \left(10^6, 10^6, \ldots, 10^6\right)^T\ , \tag{33}$$

so $\|\mathbf{S}^{-1}\mathbf{K}\mathbf{x}^*\|_2 = 1.225 \times 10^7$. Substituting these values into (32) gives $|\Delta x|_{rms} \le 3.67 \times 10^3$, a wildly pessimistic bound. Figure 3 gives a componentwise plot of the actual errors $\hat{\mathbf{x}} - \mathbf{x}^*$ with the true values of $\pm|\Delta x|_{rms} = \pm 0.302$ plotted as dashed lines.

The classical bound is hopelessly pessimistic because it does not take the random nature of the errors into account. Starting with a measured $\mathbf{b}$ and corresponding solution $\mathbf{x}$, it considers all measured vectors $\mathbf{b} + \delta\mathbf{b}$ with $\|\delta\mathbf{b}\|_2 \le \|\Delta\mathbf{b}\|_2$. These vectors define corresponding solutions $\hat{\mathbf{x}} = \mathbf{x} + \delta\mathbf{x}$, and to make the bound hold with certainty for all $\mathbf{b} + \delta\mathbf{b}$, it assumes the worst possible combination of the 121 perturbations $\delta\mathbf{b}$. When the errors are drawn randomly, the probability of such a combination is negligibly small.

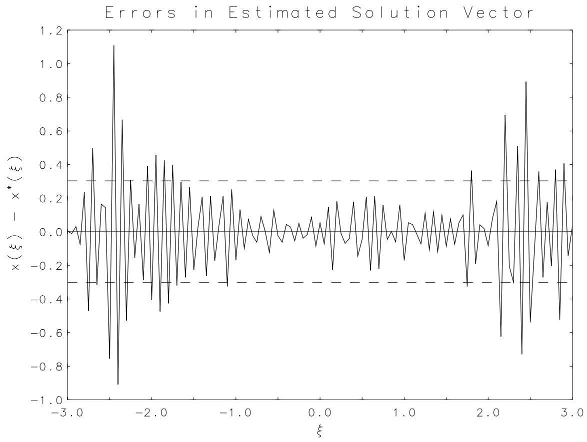

Figure 3:

Statistical Perturbation Bounds

A more reasonable bound can be obtained by considering the statistical properties of the errors. By (13),

$$(\hat{\mathbf{x}} - \mathbf{x}^*) \sim N\left[\mathbf{o}, \left(\mathbf{K}^T\mathbf{S}^{-2}\mathbf{K}\right)^{-1}\right] , \tag{34}$$

so

$$(\hat{\mathbf{x}} - \mathbf{x}^*)^T\mathbf{K}^T\mathbf{S}^{-2}\mathbf{K}(\hat{\mathbf{x}} - \mathbf{x}^*) \sim \chi^2(n) , \tag{35}$$

whence

$$\mathcal{E}\left\{(\hat{\mathbf{x}} - \mathbf{x}^*)^T\mathbf{K}^T\mathbf{S}^{-2}\mathbf{K}(\hat{\mathbf{x}} - \mathbf{x}^*)\right\} = n . \tag{36}$$

Now consider the singular value decomposition

$$\begin{gathered}\mathbf{S}^{-1}\mathbf{K} = \mathbf{U}\begin{pmatrix}\mathbf{\Sigma}\\ \mathbf{0}\end{pmatrix}\mathbf{V}^T , \quad \mathbf{\Sigma} = \mathbf{diag}(\sigma_1, \sigma_2, \ldots, \sigma_n) , \\ \mathbf{U}^T\mathbf{U} = \mathbf{I}_m , \quad \mathbf{V}^T\mathbf{V} = \mathbf{I}_n , \quad \sigma_1 \geq \sigma_2 \geq \cdots \geq \sigma_n .\end{gathered} \tag{37}$$

Substituting into (36) and simplifying gives

$$\mathcal{E}\left\{\sum_{j=1}^{n}\sigma_j^2\left[\mathbf{V}^T(\hat{\mathbf{x}} - \mathbf{x}^*)\right]_j^2\right\} = n , \tag{38}$$

and, since σ_n is the minimum singular value,

$$\sigma_n^2\mathcal{E}\left\{\sum_{j=1}^{n}\left[\mathbf{V}^T(\hat{\mathbf{x}} - \mathbf{x}^*)\right]_j^2\right\} \leq n . \tag{39}$$

Dividing through by σ_n^2 gives

$$\mathcal{E}\left\{\sum_{j=1}^{n}\left[\mathbf{V}^T(\hat{\mathbf{x}} - \mathbf{x}^*)\right]_j^2\right\} = \mathcal{E}\left\{\|\mathbf{V}^T(\hat{\mathbf{x}} - \mathbf{x}^*)\|_2^2\right\} \leq \frac{n}{\sigma_n^2} . \tag{40}$$

The two-norm is invariant with orthogonal rotations, so

$$\mathcal{E}\left\{\|\hat{\mathbf{x}} - \mathbf{x}^*\|_2^2\right\} \leq \frac{n}{\sigma_n^2} , \tag{41}$$

whence, by (31),

$$|\Delta x|_{rms} \leq \frac{1}{\sigma_n} . \tag{42}$$

This bound is computable without knowing $\mathbf{x}^*$, and *it does not depend on* $\mathrm{cond}(\mathbf{S}^{-1}\mathbf{K})$. For the test problem, $|\Delta x|_{rms} \leq 0.861$, which exceeds the true value by a factor of only 2.85.

Confidence Intervals

Both the classical and statistical perturbation analyses are rendered moot by confidence interval calculations. If $\hat{\mathbf{x}}$ is the least squares solution for the model (13), then

$$\hat{\mathbf{x}} \sim N\left[\mathbf{x}^*, \left(\mathbf{K}^T\mathbf{S}^{-2}\mathbf{K}\right)^{-1}\right] , \tag{43}$$

so the variances of the invidual $\hat{x}_j$ are given by

$$\mathrm{V}(\hat{x}_j) = \mathbf{e}_j^T\left(\mathbf{K}^T\mathbf{S}^{-2}\mathbf{K}\right)^{-1}\mathbf{e}_j^T , \quad j = 1, 2, \ldots, n , \tag{44}$$

where $\mathbf{e}_j$ is the unit vector with 1 as the jth element. For any probability α ($0 < \alpha < 1$), if κ is chosen to satisfy

$$\frac{1}{\sqrt{2\pi}}\int_{-\kappa}^{+\kappa}\exp\left(-\frac{\eta^2}{2}\right)d\eta = \alpha , \tag{45}$$

then

$$\Pr\left\{\left[\hat{x}_j - \kappa\sqrt{\mathrm{V}(\hat{x}_j)}\right] \leq x_j^* \leq \left[\hat{x}_j + \kappa\sqrt{\mathrm{V}(\hat{x}_j)}\right]\right\} = \alpha . \tag{46}$$

The κ-value for $\alpha = .95$ is $\kappa = 1.96$. Figure 4 shows the 95% confidence bounds for the test problem. The dashed line is the true solution and the jagged lines connect the upper and lower bounds for the individual $\hat{x}_i$.

If $\mathbf{S}^2 = s^2I_m$, with s unknown, then the estimate $\hat{s}^2 = (m-n)^{-1}\rho_{LS}^2$ can be used to construct confidence intervals, though the relation between κ and α will be different from (45). If the $\hat{\epsilon}$-distribution is unknown, confidence intervals can be constructed from the Chebeyshev inequality. Though wider than those for normally distributed errors, these intervals are often orders of magnitude smaller than the $\pm|\Delta x|_{rms}$ bounds from classical perturbation theory.

The keynote speaker [7] pointed out that the variance matrix for $\hat{x}_j$ was known to Gauss, and that modern least squares algorithms could easily compute it by inverting an upper triangular matrix formed in solving

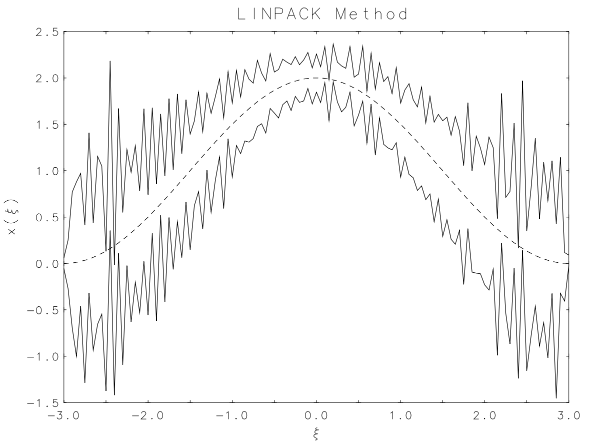

Figure 4:

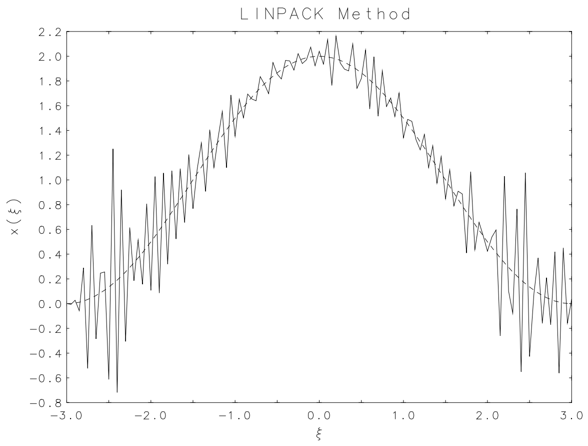

Figure 5:

for $\hat{\mathbf{x}}$. Unfortunately, the least squares subroutines in the widely used LINPACK [2] and LAPACK [1] collections do not return confidence intervals, or even the variance matrix. The LINPACK manual describes how to compute variances from a reduced matrix returned by subroutine SQRDC, but the LAPACK manual is silent on the subject, and neither mentions confidence intervals, concentrating instead on the classical perturbation bounds. Secondary sources, which use these collections, have continued this preoccupation with what are essentially useless bounds. They also continue to propagate misinformation about the condition number. For example, the textbook of Kahaner, et. al [4] states that:

> One useful interpretation of the condition number is that its logarithm approximates the number of digits which will be lost while solving $Ax = b$. Thus if $\text{cond}(A) = 10^5$ and if machine epsilon is 10^{-8}, then the best we can expect is that the solution will be accurate to about three digits.

The estimate in Figure 2 was calculated in double precision with $\epsilon_{\text{mach}} = 2.22 \times 10^{-16}$, and since $\text{cond}(\mathbf{S}^{-1}\mathbf{K}) = 2.92 \times 10^9$, the above reasoning would indicate that the computed $\hat{\mathbf{x}}$ is accurate to 6 digits. But consider the same calculation in single precision with $\epsilon_{\text{mach}} = 1.19 \times 10^{-7}$ and $\text{cond}(\mathbf{S}^{-1}\mathbf{K}) = 2.93 \times 10^9$. According to the conventional wisdom, a computed estimate should not contain any digits of accuracy. The actual single precision estimate is shown in Figure 5. The slight differences from the double precision estimate are difficult to see by comparing the two plots. The rms average difference between the two estimates is 0.0033 which is almost 100 time smaller than the $|\Delta x|_{rms}$ for either estimate, so in practice, either estimate would serve equally well. Clearly the condition number is not always a good indicator of the accuracy of the estimate.

Acknowledgements

I would like to thank Dr. W.R. Burrus for pointing out that the classical perturbation bounds are usually too pessimistic to be useful and Dr. R.F. Boisvert for many useful suggestions for improving the manuscript.

References

[1] Anderson, E., Bai, Z., Bischof, C., Demmel, J., Dongarra, J., Du Croz, J., Greenbaum, A., Hammarling, S., McKenney, A., Ostrouchov, S., and Sorensen, D. (1992) *LAPACK User's Guide*, SIAM, Philadelphia.

[2] Dongarra, J.J., Moler, C.B., Bunch, J.R., and Stewart, G.W. (1979) *LINPACK Users' Guide*, SIAM, Philadelphia, Chapt. 9.

[3] Golub, G.H. and Van Loan, C.F. (1989) *Matrix Computations*, The Johns Hopkins University Press, Baltimore, Chapt. 5.

[4] Kahaner, D., Moler, C., and Nash, S. (1989) *Numerical Methods and Software*, Prentice Hall, Englewood Cliffs, Chapt. 3.

[5] Phillips, D.L. (1962) *J. Assoc. Comput. Mach.*, **9**, pp. 84-97.

[6] Stewart, G.W. (1973) *Introduction to Matrix Computations*, Academic Press, New York, Chapt. 5.

[7] Stewart, G.W. (1994) *This Volume.*