

Residual periodograms for choosing regularization parameters for ill-posed problems

Bert W Rust¹ and Dianne P O'Leary^{1,2}

¹ Mathematical and Computational Sciences Division, National Institute of Standards and Technology, Gaithersburg, MD 20899, USA

² Computer Science Department and Institute for Advanced Computer Studies, University of Maryland, College Park, MD 20742, USA

E-mail: bert.rust@nist.gov and oleary@cs.umd.edu

Received 28 June 2007, in final form 6 October 2007

Published 23 May 2008

Online at stacks.iop.org/IP/24/034005

Abstract

Consider an ill-posed problem transformed if necessary so that the errors in the data are independent identically normally distributed with mean zero and variance 1. We survey regularization and parameter selection from a linear algebra and statistics viewpoint and compare the statistical distributions of regularized estimates of the solution and the residual. We discuss methods for choosing a regularization parameter in order to assure that the residual for the model is statistically plausible. Ideally, as proposed by Rust (1998 *Tech. Rep. NISTIR 6131*, 2000 *Comput. Sci. Stat.* **32** 333–47), the results of candidate parameter choices should be evaluated by plotting the resulting residual along with its periodogram and its cumulative periodogram, but sometimes an automated choice is needed. We evaluate a method for choosing the regularization parameter that makes the residuals as close as possible to white noise, using a diagnostic test based on the periodogram. We compare this method with standard techniques such as the discrepancy principle, the L-curve and generalized cross validation, showing that it performs better on two new test problems as well as a variety of standard problems.

1. Introduction

Systems of first kind integral equations,

$$y_i \equiv y(t_i) = \int_a^b K(t_i, \xi)x(\xi) d\xi + \epsilon_i, \quad i = 1, 2, \dots, m, \quad (1.1)$$

are routinely used to model instrument distortions in measuring an unknown function $x(\xi)$. The $K_i(\xi) \equiv K(t_i, \xi)$ are known (previously measured or calculated) response functions of

the instrument, the y_i are measurements made on a discrete mesh t_1, t_2, \dots, t_m , and the ϵ_i are random, zero-mean measuring errors. Discretizing gives a linear regression model

$$\mathbf{y} = \mathbf{K}\mathbf{x}^* + \boldsymbol{\epsilon}, \quad (1.2)$$

where \mathbf{y} is the m -vector of measurements, \mathbf{K} is a known $m \times n$ matrix, with $m \geq n$, and \mathbf{x}^* is an unknown n -vector whose components are either discrete values of $x(\xi)$ on some mesh $\xi_1, \xi_2, \dots, \xi_n$, or are the unknown coefficients in a truncated expansion for $x(\xi)$. An example of such a discretization is given in appendix A. The vector $\boldsymbol{\epsilon}$ is an m -vector of random measuring errors satisfying

$$\mathcal{E}(\boldsymbol{\epsilon}) = \mathbf{0}, \quad \mathcal{E}(\boldsymbol{\epsilon}\boldsymbol{\epsilon}^T) = \mathbf{S}^2, \quad (1.3)$$

where \mathcal{E} is the expectation operator, $\mathbf{0}$ is the zero vector and \mathbf{S}^2 is the positive definite variance matrix³ for $\boldsymbol{\epsilon}$. When the measurement errors are statistically uncorrelated, the variance matrix is diagonal:

$$\mathbf{S}^2 = \text{diag}(s_1^2, s_2^2, \dots, s_m^2), \quad (1.4)$$

where s_1, s_2, \dots, s_m are the standard deviations of the errors. When the errors are correlated, the model can be transformed to have a diagonal variance matrix by premultiplying (1.2) by the inverse of the lower triangular Cholesky factor of \mathbf{S}^2 .

Estimates of \mathbf{S} are often readily available along with the data values \mathbf{y} , since good experimenters routinely provide estimates of the error as ± 1 -sigma error bars on the plotted points. The variances can be estimated as the squares of the half-lengths of the bars}. An analyst who fails to use this information implicitly assumes that $\mathbf{S}^2 = s^2 \mathbf{I}_m$, where \mathbf{I}_m is the m th-order identity matrix and s is an unknown scalar that can be, but usually is not, estimated from the sum of squared residuals for the least squares solution. Using all available information on the variances can greatly improve estimates of \mathbf{x} .

In the following it will be assumed that \mathbf{S} is a known matrix and that the errors are samples from a multivariate normal distribution, i.e., that $\boldsymbol{\epsilon} \sim N(\mathbf{0}, \mathbf{S}^2)$. In section 2, these assumptions will be used to rescale the problem so that the scaled errors $\boldsymbol{\eta} = \mathbf{S}^{-1}\boldsymbol{\epsilon} \sim N(\mathbf{0}, \mathbf{I}_m)$ and to derive a statistical diagnostic for estimates of the solution vector. In section 3 we use this scaling on a variant of the well-known Phillips problem [23]; the ordinary least squares estimate is calculated and found to be unsatisfactory. In section 4 we discuss the families of solutions formed from Tikhonov estimation and from the truncated singular value decomposition (TSVD), among which an acceptable estimate of the solution can often be found, and derive their statistical properties. Section 5 discusses rules for choosing among the estimates. In section 6 we illustrate these ideas on a test problem. The new diagnostics can be used to specify high-quality Tikhonov and TSVD estimates, but in both cases, the L-curve and minimum generalized cross validation (GCV) criteria give unacceptable results. In section 7 we discuss how the parameter choice method can be automated, and we present results on a variety of standard test problems in section 8. In section 9 the methods are successfully applied to real-world measurements of the energy spectrum of neutrons produced by a certain nuclear reaction. Finally, section 10 gives a brief discussion of algorithmic considerations and of how the new method can be extended when the knowledge of the measurement errors is not as complete as might be desired.

In some important ill-posed problems, extra information is available about the solution; for example, we might know that $x(\xi)$ is non-negative or monotonic. Including such constraints can be important in achieving a good solution using techniques such as penalized maximum likelihood, Bayesian methods, maximum entropy, etc (see for example [2, 4, 19, 24, 29]), but

³ This matrix is often called the *covariance matrix*.

we do not consider such constraints here. Neither do we consider the important question of robustness of our methods when our assumptions on the distribution of the error are violated.

This paper is meant to be expository, providing an overview of methods for choosing regularization parameters. At the same time, it contains some original contributions. We compare the statistical distributions of regularized estimates of the solution and the residual for discretized ill-posed problems; we highlight a particularly promising and underutilized set of statistical techniques that we have found useful in judging the quality of a solution; we propose a way to automate the evaluation when necessary, and we present two new test problems⁴ useful for evaluating numerical methods. Some of this material is taken from [26].

2. Properties of the model's residual

Our linear regression model can be written as

$$\mathbf{y} = \mathbf{K}\mathbf{x}^* + \boldsymbol{\epsilon}, \quad \boldsymbol{\epsilon} \sim N(\mathbf{0}, \mathbf{S}^2), \quad (2.5)$$

as stated in (1.2), but it is advantageous to scale it with the matrix \mathbf{S}^{-1} . Let

$$\mathbf{b} \equiv \mathbf{S}^{-1}\mathbf{y}, \quad \mathbf{A} \equiv \mathbf{S}^{-1}\mathbf{K}, \quad \boldsymbol{\eta} \equiv \mathbf{S}^{-1}\boldsymbol{\epsilon}. \quad (2.6)$$

Note that by a standard theorem of multivariate statistics [1, theorem 2.4.4], $\boldsymbol{\eta} \sim N(\mathbf{S}^{-1}\mathbf{0}, \mathbf{S}^{-1}\mathbf{S}^2[\mathbf{S}^{-1}]^T)$, so the scaled model can be written as

$$\mathbf{b} = \mathbf{A}\mathbf{x}^* + \boldsymbol{\eta}, \quad \boldsymbol{\eta} \sim N(\mathbf{0}, \mathbf{I}_m). \quad (2.7)$$

As a consequence, another standard theorem [17, p 140] gives $\|\boldsymbol{\eta}\|^2 \sim \chi^2(m)$, where $\chi^2(m)$ is the Chi-squared distribution with m degrees of freedom. The advantage of our scaling is that we now can see how a reasonable residual to the model should behave. Let $\tilde{\mathbf{x}}$ be an estimate of \mathbf{x}^* and let

$$\tilde{\mathbf{r}} = \mathbf{b} - \mathbf{A}\tilde{\mathbf{x}} \quad (2.8)$$

be the corresponding residual vector. Since the regression model can also be written as

$$\boldsymbol{\eta} = \mathbf{b} - \mathbf{A}\mathbf{x}^*, \quad (2.9)$$

it is clear that $\tilde{\mathbf{x}}$ is *acceptable only if $\tilde{\mathbf{r}}$ is a plausible sample from the distribution from which $\boldsymbol{\eta}$ is drawn*. Since

$$\mathcal{E}\{\|\boldsymbol{\eta}\|^2\} = m, \quad \text{Var}\{\|\boldsymbol{\eta}\|^2\} = 2m, \quad (2.10)$$

these two quantities provide rough bounds for the $\|\tilde{\mathbf{r}}\|^2$ that might be expected from a reasonable estimate of \mathbf{x}^* : an estimate that gives

$$m - \sqrt{2m} \leq \|\mathbf{b} - \mathbf{A}\tilde{\mathbf{x}}\|^2 \leq m + \sqrt{2m} \quad (2.11)$$

would be reasonable, but any $\tilde{\mathbf{x}}$ with $\|\tilde{\mathbf{r}}\|^2$ outside the interval $[m - 2\sqrt{2m}, m + 2\sqrt{2m}]$ would be suspect. These indicators can be sharpened and quantified by using percentiles of the cumulative distribution function for $\chi^2(m)$.

We will see in section 5 how to use this information in choosing among a family of possible *regularized* estimates for \mathbf{x}^* .

⁴ The data for the Burrus problem in section 9 and a MATLAB program for the Phillips variant defined in the appendix are available at www.cs.umd.edu/users/oleary/bwr.

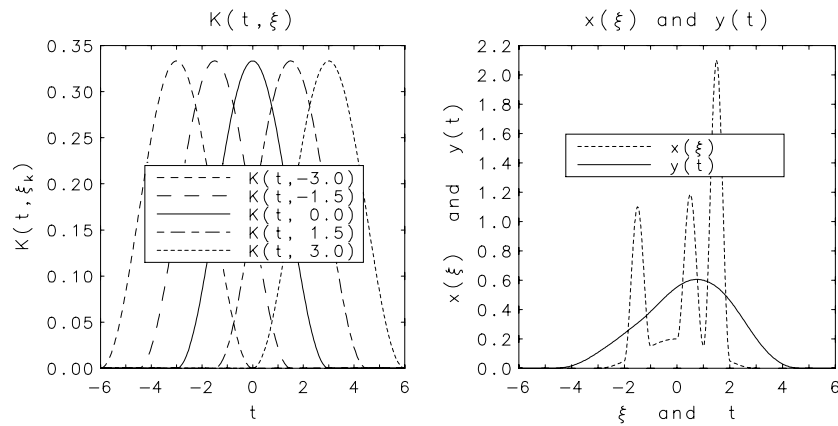


Figure 1. The modified Phillips problem. Left: the kernel $K(t, \xi)$ for various values of ξ , from (A.2) in the Appendix. Right: the true solution $x(\xi)$ from (A.3) and the function $y(t)$ from (A.1).

3. Linear regression

By transforming the matrix model from \mathbf{K} to \mathbf{A} , we equalize the contributions of errors in the data, but we have not mitigated the sensitivity of the model to these errors. If the problem were not ill posed, we would assume that $\text{rank}(\mathbf{A}) = n$ and seek the minimum variance, unbiased estimator by solving

$$r_{\min}^2 = \min_{\mathbf{x} \in \mathbb{R}^n} \{(\mathbf{b} - \mathbf{Ax})^T (\mathbf{b} - \mathbf{Ax})\}. \quad (3.12)$$

The linear regression estimate

$$\hat{\mathbf{x}} = (\mathbf{A}^T \mathbf{A})^{-1} \mathbf{A}^T \mathbf{b} \quad (3.13)$$

is the *best linear unbiased estimate* of \mathbf{x}^* [1]. To compute its variance, we use the fact that for any matrix \mathbf{C} with m columns, if $\mathbf{b} \sim N(\mathbf{Ax}^*, \mathbf{I}_m)$, then [1, theorem 2.4.1]

$$\mathbf{Cb} \sim N(\mathbf{CAx}^*, \mathbf{CC}^T), \quad (3.14)$$

so the variance of $\hat{\mathbf{x}}$ is $(\mathbf{A}^T \mathbf{A})^{-1}$. A similar computation shows that the distribution of the corresponding residual has mean $\mathbf{0}$ and variance

$$\mathbf{I}_m - \mathbf{A}(\mathbf{A}^T \mathbf{A})^{-1} \mathbf{A}^T. \quad (3.15)$$

It is well known [28, chapter 1; 32, chapter 6; 35, chapter 2] that for ill-posed problems, the elements of $\hat{\mathbf{x}}$ are pathologically sensitive to small variations in the elements of \mathbf{b} , so despite the desirable statistical properties of the estimator $\hat{\mathbf{x}}$, the measuring errors generally make it totally unphysical and wildly oscillating.

We illustrate this on a useful test problem that shares many characteristics of real instrument correction problems. This problem can be obtained by discretizing a variant of the Phillips equation [23] and adding random errors to the discrete y_i . The problem and a discretization with $m = 300$ and $n = 241$ are described in detail in appendix A. The functions $y(t)$ and $x(\xi)$ are plotted on the right-hand side in figure 1, and on the left-hand side we see the functions $K(t, \xi_j)$ for the discrete values $\xi_j = -3.0, -1.5, 0, 1.5, 3.0$. All of the 241 $K(t, \xi_k)$ have the same shape and subtend unit area. The 300 standard deviations s_i for equation (A.24) in the Appendix range in value from 3.49×10^{-11} to 7.78×10^{-6} , so the errors in the y_i are much smaller than the thickness of the curve in the plot of $y(t)$. But these small errors produce large oscillations in the least squares estimate of the solution vector. This is shown in the

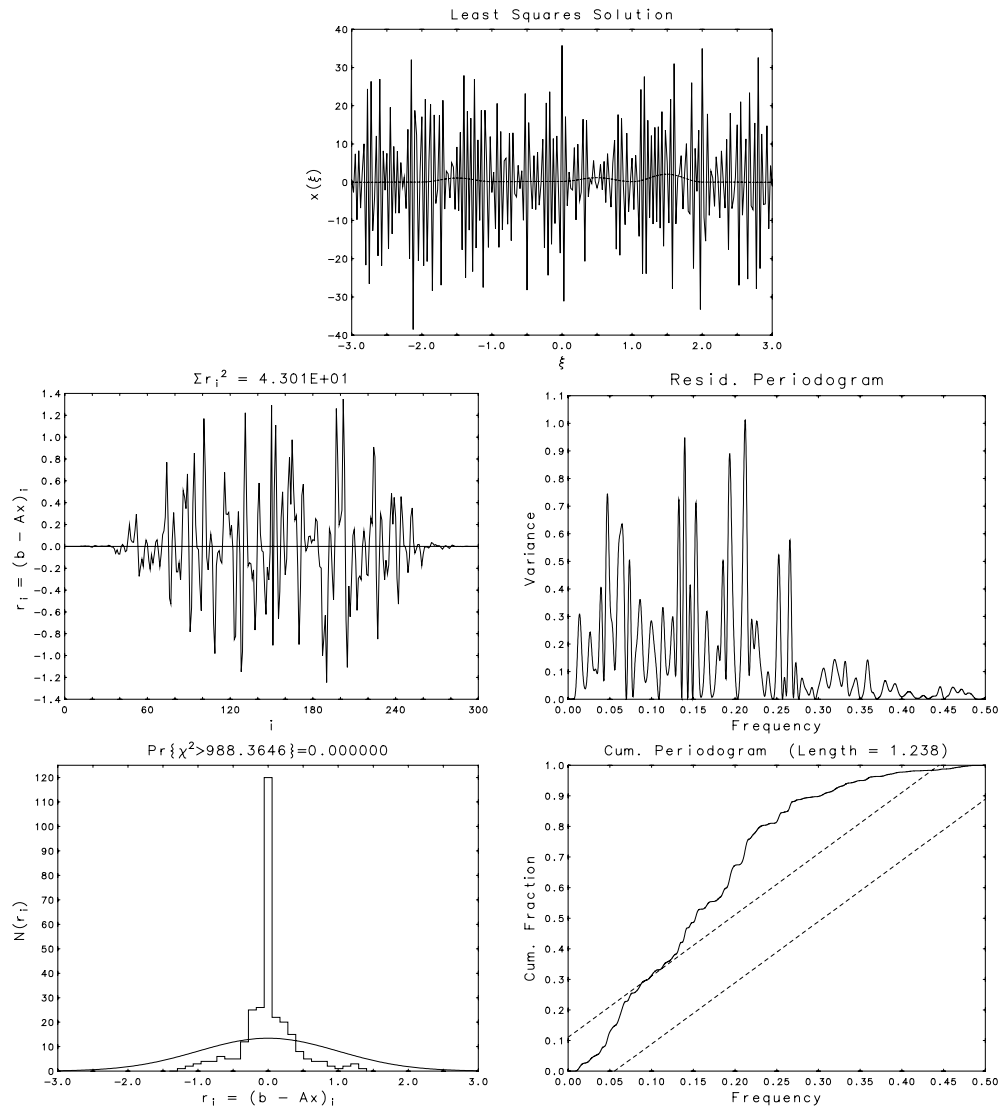


Figure 2. Least squares estimate (top) and diagnostics for the modified Phillips problem. On the left are the residual and its histogram. On the right are the residual periodogram and the cumulative periodogram, discussed in section 5.2.

top plot of figure 2 where the barely discernible dashed curve is the true solution \mathbf{x}^* , and the solid curve oscillating wildly about it is the least squares estimate $\hat{\mathbf{x}}$. The magnitudes of these oscillations are roughly 10^7 times greater than the largest random errors in the y_i . Such large amplifications of the measuring errors are typical for ill-posed problems.

Even without having information about the true solution, the norm of the residual tells us that the computed solution is not useful: for the modified Phillips problem, the linear regression estimate gives

$$r_{\min}^2 = \min_{\mathbf{x} \in R^n} \|\mathbf{b} - \mathbf{Ax}\|^2 = \|\mathbf{b} - \mathbf{A}\hat{\mathbf{x}}\|^2 = 43.01. \quad (3.16)$$

By (2.10), $\mathcal{E}\{\|\mathbf{b} - \mathbf{A}\mathbf{x}^*\|^2\} = 300$, with standard deviation $\sqrt{600} = 24.49$, so r_{\min}^2 is more than 10 standard deviations smaller than the expected value.

To improve our estimate, we need a family of *regularized solution estimates* that are not so sensitive to errors, and some means for choosing among those estimates.

4. Regularized solution estimates

Insight into the failure of the least squares method is obtained by use of the *singular value decomposition* (SVD) of \mathbf{A} :

$$\mathbf{A} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T = \mathbf{U} \begin{bmatrix} \mathbf{\Sigma}_1 \\ \mathbf{O} \end{bmatrix} \mathbf{V}^T, \quad \mathbf{\Sigma}_1 = \mathbf{diag}(\sigma_1, \sigma_2, \dots, \sigma_n). \quad (4.17)$$

Here $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_n \geq 0$, and

$$\mathbf{U}^T \mathbf{U} = \mathbf{I}_m = \mathbf{U} \mathbf{U}^T, \quad \mathbf{V}^T \mathbf{V} = \mathbf{I}_n = \mathbf{V} \mathbf{V}^T. \quad (4.18)$$

If the $m \times m$ matrix \mathbf{U} is partitioned

$$\mathbf{U} = [\mathbf{U}_1, \mathbf{U}_2], \quad (4.19)$$

with \mathbf{U}_1 being an $m \times n$ submatrix, then it can be shown by substituting into (3.13) [8, section 5.5.3] that the least squares solution satisfies

$$\mathbf{V}^T \hat{\mathbf{x}} = \mathbf{\Sigma}_1^{-1} \mathbf{U}_1^T \mathbf{b}, \quad (4.20)$$

and that

$$r_{\min}^2 = \|\mathbf{b} - \mathbf{A}\hat{\mathbf{x}}\|^2 = \|\mathbf{U}_2^T \mathbf{b}\|^2. \quad (4.21)$$

These last two equations can also be written as

$$(\mathbf{V}^T \hat{\mathbf{x}})_i = \frac{(\mathbf{U}^T \mathbf{b})_i}{\sigma_i}, \quad i = 1, 2, \dots, n, \quad (4.22)$$

and

$$r_{\min}^2 = \|\mathbf{b} - \mathbf{A}\hat{\mathbf{x}}\|^2 = \sum_{i=n+1}^m (\mathbf{U}^T \mathbf{b})_i^2. \quad (4.23)$$

Using the SVD of \mathbf{A} , we can compute $\mathbf{A} = \mathbf{U}_1 \mathbf{\Sigma}_1 \mathbf{V}^T$ and

$$\mathbf{A}(\mathbf{A}^T \mathbf{A})^{-1} \mathbf{A}^T = \mathbf{U}_1 \mathbf{U}_1^T. \quad (4.24)$$

This analysis leads to two families of regularized solutions: truncated SVD (TSVD) and Tikhonov regularization.

4.1. The truncated SVD family of regularized solutions

For the modified Phillips problem, the value of r_{\min}^2 defined by (3.16) is too small by ten standard deviations, which suggest that some of the $(\mathbf{U}^T \mathbf{b})_i$ values in the sequence (4.22) more properly belong to the sum in (4.23). This reasoning leads to the idea of truncating the decomposition. This is accomplished by choosing a value $p < n$ and replacing $\mathbf{\Sigma}_1$ by a truncated matrix

$$\mathbf{\Sigma}_{tr} = \mathbf{diag}(\sigma_1, \dots, \sigma_p, 0, \dots, 0) \quad (4.25)$$

whose pseudo-inverse is

$$\mathbf{\Sigma}_{tr}^\dagger = \mathbf{diag}\left(\frac{1}{\sigma_1}, \dots, \frac{1}{\sigma_p}, 0, \dots, 0\right). \quad (4.26)$$




Figure 3. Singular values (squares) and first n elements of $|\mathbf{U}^T \mathbf{b}|$ (circles) for the modified Phillips problem.

Then (4.22) and (4.23) are replaced by

$$(\mathbf{V}^T \tilde{\mathbf{x}})_i = \begin{cases} \frac{1}{\sigma_i} (\mathbf{U}^T \mathbf{b})_i, & i = 1, 2, \dots, p, \\ 0, & i = p+1, \dots, n. \end{cases} \quad (4.27)$$

From this, we can compute

$$\|\mathbf{b} - \mathbf{A}\tilde{\mathbf{x}}\|^2 = \sum_{i=p+1}^m (\mathbf{U}^T \mathbf{b})_i^2. \quad (4.28)$$

The above approach was first suggested by Golub and Kahan [6] who noted its similarity to the theoretical treatment given by Smithies [27, chapter 8] for the singular functions and singular values of first kind integral equations. One of the first to use it was Richard Hanson [14] who suggested that the threshold should be the smallest integer p such that

$$\sum_{i=p+1}^m (\mathbf{U}^T \mathbf{b})_i^2 < m. \quad (4.29)$$

In view of (2.10), this seems a very sensible choice. Another choice [7] is to seek a clear gap in the distribution of the σ_i and to zero all those on the low side. Unfortunately, most ill-posed problems have no gap.

As an example, we see no gap in the plot of the singular values for the modified Phillips problem, marked with squares in figure 3, which also shows the elements of $|\mathbf{U}^T \mathbf{b}|$ plotted as connected circles. The largest and smallest singular values are

$$\sigma_1(\mathbf{A}) = 2.882 \times 10^5, \quad \sigma_{241}(\mathbf{A}) = 2.622 \times 10^{-2}. \quad (4.30)$$

with ratio (condition number) $\text{cond}(\mathbf{A}) = 1.099 \times 10^7$. The relative accuracy is $\epsilon_{\text{mach}} = 2.22 \times 10^{-16}$, so there is no reason to assume that the numerical rank of \mathbf{A} is less than n . Yet the problem obviously needs some truncation to prevent the estimate from capturing variance that properly belongs to the residuals.

We note that the TSVD solution estimate (4.27) can also be obtained by zeroing components of $\mathbf{U}^T \mathbf{b}$ corresponding to small singular values rather than zeroing the singular values of \mathbf{A} , as noted by Golub and Kahan [6].

4.2. The Tikhonov family of regularized solutions

An alternative to setting small singular values to zero is to instead increase them so that they lead to a smaller contribution to the solution estimate; see (4.22). Although this strategy is generally attributed to Tikhonov [30], it was also pioneered by Phillips [23] and Twomey [31]. A statistical variant, called *ridge regression*, was also independently developed by Hoerl and Kennard [15, 16]. In the Tikhonov method, we add a small increment to each singular value. For the large singular values this has little effect, but for the small ones it can greatly diminish the noise contributed by the components of $\mathbf{U}_1^T \mathbf{b}$ corresponding to those values.

The Tikhonov regularization estimate is given by

$$\tilde{\mathbf{x}}_\lambda = (\mathbf{A}^T \mathbf{A} + \lambda^2 \mathbf{I}_n)^{-1} \mathbf{A}^T \mathbf{b}, \quad (4.31)$$

where λ is a parameter chosen to balance the competing demands of fidelity to the measurements and insensitivity to measurement errors. This interpretation of λ comes from the fact that the vector $\tilde{\mathbf{x}}_\lambda$ is the solution to the minimization problem

$$\min_{\mathbf{x}} \|\mathbf{b} - \mathbf{A}\mathbf{x}\|^2 + \lambda^2 \|\mathbf{x}\|^2. \quad (4.32)$$

Using the SVD, it can be computed as

$$(\mathbf{V}^T \tilde{\mathbf{x}})_i = \frac{\sigma_i}{\sigma_i^2 + \lambda^2} (\mathbf{U}^T \mathbf{b})_i, \quad i = 1, 2, \dots, n, \quad (4.33)$$

which can be compared to the TSVD value (4.27). The corresponding residual norm-squared is

$$\tilde{r}_{\min}^2 = \|\mathbf{b} - \mathbf{A}\tilde{\mathbf{x}}\|^2 = \sum_{i=1}^n \left(\frac{\lambda^2}{\sigma_i^2 + \lambda^2} \right)^2 (\mathbf{U}^T \mathbf{b})_i^2 + \sum_{i=n+1}^m (\mathbf{U}^T \mathbf{b})_i^2, \quad (4.34)$$

which compares to the TSVD value (4.28).

4.3. The statistical distributions of our estimators

In this section we use the SVD of the matrix \mathbf{A} to derive formulas for the means and variances of the estimators obtained from least squares and from our regularized algorithms. We will repeatedly make use of several useful properties of the SVD. First, since the columns of \mathbf{U} and \mathbf{V} form orthonormal bases, we have the properties $\mathbf{U}^T \mathbf{U} = \mathbf{I}_m$ and $\mathbf{V}^T \mathbf{V} = \mathbf{I}_n$. Second, since Σ is an $m \times n$ matrix that is nonzero only on the main diagonal, we will call its top $n \times n$ block Σ_1 and denote the $n \times m$ pseudo-inverse of Σ by

$$\Sigma^\dagger = [\Sigma_1^\dagger \quad \mathbf{0}], \quad (4.35)$$

where the pseudo-inverse of a square diagonal matrix is formed by replacing the nonzero diagonal entries by their inverses. To make expressions simpler we will assume that \mathbf{A} and therefore Σ_1 are full rank so that $\Sigma_1^\dagger = \Sigma_1^{-1}$ and $\Sigma^\dagger \Sigma = \mathbf{I}_n$. We will also make use of the partitioning of \mathbf{U} in equation (4.19), and the fact that diagonal matrices commute. We have three kinds of solution estimates for problem (1.2): the least squares estimate, the TSVD estimate and the Tikhonov estimate. Referring to equations (4.22), (4.27) and (4.33), we see that all of these estimates have a similar form. In fact, for all of them,

$$\mathbf{V}^T \mathbf{x} = \mathbf{F} \Sigma^\dagger \mathbf{U}^T \mathbf{b}, \quad (4.36)$$

where \mathbf{F} is an $n \times n$ diagonal matrix specific to each method. The entries of \mathbf{F} are called *filter factors* [11]. For the ordinary least squares estimate, \mathbf{F} is the identity matrix. For Tikhonov regularization, the j th diagonal element of \mathbf{F} is

$$f_j = \frac{\sigma_j^2}{\sigma_j^2 + \lambda^2}. \quad (4.37)$$

For TSVD, the first p diagonal entries of \mathbf{F} are 1 and the others are zero.

So in each case we can write the solution estimate as

$$\mathbf{x} = \mathbf{V}\mathbf{F}\mathbf{\Sigma}^\dagger\mathbf{U}^T\mathbf{b} \equiv \mathbf{C}\mathbf{b}. \quad (4.38)$$

From this we see that under the assumption that $\mathbf{b} \sim N(\mathbf{A}\mathbf{x}^*, \mathbf{I}_m)$, \mathbf{x} has a normal distribution. We can compute the mean and variance of the distribution by using (3.14). We see, by setting $\mathbf{C} = \mathbf{V}\mathbf{F}\mathbf{\Sigma}^\dagger\mathbf{U}^T$ and $\mathbf{b}^* = \mathbf{A}\mathbf{x}^*$, that the mean of \mathbf{x} is

$$\mathbf{C}\mathbf{b}^* = \mathbf{V}\mathbf{F}\mathbf{\Sigma}^\dagger\mathbf{U}^T\mathbf{A}\mathbf{x}^* \quad (4.39)$$

$$= \mathbf{V}\mathbf{F}\mathbf{\Sigma}^\dagger\mathbf{U}^T\mathbf{U}\mathbf{\Sigma}\mathbf{V}^T\mathbf{x}^* \quad (4.40)$$

$$= \mathbf{V}\mathbf{F}\mathbf{V}^T\mathbf{x}^* \quad (4.41)$$

$$= \mathbf{x}^* - \mathbf{V}(\mathbf{I}_n - \mathbf{F})\mathbf{V}^T\mathbf{x}^*, \quad (4.42)$$

and the variance of \mathbf{x} is

$$\mathbf{C}\mathbf{C}^T = \mathbf{V}\mathbf{F}\mathbf{\Sigma}^\dagger\mathbf{U}^T\mathbf{U}(\mathbf{\Sigma}^\dagger)^T\mathbf{F}\mathbf{V}^T \quad (4.43)$$

$$= \mathbf{V}\mathbf{F}^2(\mathbf{\Sigma}_1^2)^{-1}\mathbf{V}^T. \quad (4.44)$$

We see from these expressions that \mathbf{x} is a biased estimator if $\mathbf{F} \neq \mathbf{I}_n$, i.e., if we use a Tikhonov estimate with $\lambda \neq 0$ or a TSVD estimate with $p < n$. We also see that the variance of the estimator decreases as the filter factors decrease. Therefore the variance decreases as the Tikhonov parameter λ increases and as the TSVD parameter p decreases.

We can make a similar computation for the residual. Let $\check{\mathbf{F}}$ denote the $m \times m$ matrix that is zero except for the matrix \mathbf{F} in its upper left block. In terms of \mathbf{F} and the SVD of \mathbf{A} , the residual can be expressed as

$$\mathbf{r} = \mathbf{b} - \mathbf{A}\mathbf{x} \quad (4.45)$$

$$= \mathbf{b} - (\mathbf{U}\mathbf{\Sigma}\mathbf{V}^T)\mathbf{V}\mathbf{F}\mathbf{\Sigma}^\dagger\mathbf{U}^T\mathbf{b} \quad (4.46)$$

$$= \mathbf{U}(\mathbf{I}_m - \check{\mathbf{F}})\mathbf{U}^T\mathbf{b} \quad (4.47)$$

$$= \mathbf{U}_2\mathbf{U}_2^T\mathbf{b} + \mathbf{U}_1(\mathbf{I}_n - \mathbf{F})\mathbf{U}_1^T\mathbf{b}. \quad (4.48)$$

Note that $\mathbf{U}_2\mathbf{U}_2^T\mathbf{b}$ is the residual for least squares. Since $\mathbf{r} = \mathbf{C}\mathbf{b}$, where we redefine $\mathbf{C} = \mathbf{U}(\mathbf{I}_m - \check{\mathbf{F}})\mathbf{U}^T$, we see that \mathbf{r} is normally distributed with mean

$$\mathbf{C}\mathbf{A}\mathbf{x}^* = \mathbf{U}(\mathbf{I}_m - \check{\mathbf{F}})\mathbf{U}^T\mathbf{U}\mathbf{\Sigma}\mathbf{V}^T\mathbf{x}^* \quad (4.49)$$

$$= \mathbf{U}(\mathbf{I}_m - \check{\mathbf{F}})\mathbf{\Sigma}\mathbf{V}^T\mathbf{x}^* \quad (4.50)$$

$$= \mathbf{A}\mathbf{x}^* - \mathbf{U}_1\mathbf{F}\mathbf{\Sigma}_1\mathbf{V}^T\mathbf{x}^* \quad (4.51)$$

$$= \mathbf{U}_1\mathbf{\Sigma}_1\mathbf{V}^T\mathbf{x}^* - \mathbf{U}_1\mathbf{F}\mathbf{\Sigma}_1\mathbf{V}^T\mathbf{x}^* \quad (4.52)$$

$$= \mathbf{U}_1(\mathbf{I}_n - \mathbf{F})\mathbf{\Sigma}_1\mathbf{V}^T\mathbf{x}^*. \quad (4.53)$$

Table 1. Estimators, means, and variances for the methods in terms of the filter factor matrix \mathbf{F} and the SVD of \mathbf{A} . Each of the estimators is normally distributed.

Estimator	Mean	Variance
$\mathbf{x} = \mathbf{V}\mathbf{F}\Sigma_1^\dagger\mathbf{U}^T\mathbf{b}$	$\mathbf{x}^* - \mathbf{V}(\mathbf{I}_n - \mathbf{F})\mathbf{V}^T\mathbf{x}^*$	$\mathbf{V}(\mathbf{F}\Sigma_1^{-1})^2\mathbf{V}^T$
$\mathbf{r} = \mathbf{U}_2\mathbf{U}_2^T\mathbf{b} + \mathbf{U}_1(\mathbf{I}_n - \mathbf{F})\mathbf{U}_1^T\mathbf{b}$	$\mathbf{U}_1(\mathbf{I}_n - \mathbf{F})\Sigma_1\mathbf{V}^T\mathbf{x}^*$	$\mathbf{U}_2\mathbf{U}_2^T + \mathbf{U}_1(\mathbf{I}_n - \mathbf{F})^2\mathbf{U}_1^T$

Table 2. Estimators, means, variances, and filter factors for the three methods. Each of the estimators is normally distributed. The $n \times n$ matrix \mathbf{F} is diagonal with entries specified by the filter factors f_j , $j = 1, \dots, n$. The matrix Σ_{tr} is an $n \times n$ diagonal matrix containing the first p singular values and \mathbf{U}_{tr} , which is $m \times p$, contains the first p columns of \mathbf{U} .

	Least squares	Tikhonov	TSVD
Filter factors	$f_j = 1$	$f_j = \sigma_j^2 / (\sigma_j^2 + \lambda^2)$	$f_j = 1$ ($j \leq p$) or 0 ($j > p$)
\mathbf{x} -estimator	$\hat{\mathbf{x}} = (\mathbf{A}^T\mathbf{A})^{-1}\mathbf{A}^T\mathbf{b}$	$\mathbf{x}_\lambda = (\mathbf{A}^T\mathbf{A} + \lambda^2\mathbf{I}_n)^{-1}\mathbf{A}^T\mathbf{b}$	$\tilde{\mathbf{x}} = (\mathbf{A}^T\mathbf{A})^{-1}\mathbf{A}^T\mathbf{U}_{tr}\mathbf{U}_{tr}^T\mathbf{b}$
\mathbf{C}_x	$(\mathbf{A}^T\mathbf{A})^{-1}\mathbf{A}^T$	$(\mathbf{A}^T\mathbf{A} + \lambda^2\mathbf{I}_n)^{-1}\mathbf{A}^T$	$(\mathbf{A}^T\mathbf{A})^{-1}\mathbf{A}^T\mathbf{U}_{tr}\mathbf{U}_{tr}^T$
Mean	\mathbf{x}^*	$(\mathbf{A}^T\mathbf{A} + \lambda^2\mathbf{I}_n)^{-1}\mathbf{A}^T\mathbf{A}\mathbf{x}^*$	$(\mathbf{A}^T\mathbf{A})^{-1}\mathbf{A}^T\mathbf{U}_{tr}\mathbf{U}_{tr}^T\mathbf{A}\mathbf{x}^*$
Variance	$(\mathbf{A}^T\mathbf{A})^{-1}$ $= \mathbf{V}\Sigma_1^{-2}\mathbf{V}^T$	$(\mathbf{A}^T\mathbf{A} + \lambda^2\mathbf{I}_n)^{-2}\mathbf{A}^T\mathbf{A}$ $= \mathbf{V}\mathbf{F}^2\Sigma_1^{-2}\mathbf{V}^T$	$(\mathbf{A}^T\mathbf{A})^{-1}\mathbf{A}^T\mathbf{U}_{tr}\mathbf{U}_{tr}^T\mathbf{A}(\mathbf{A}^T\mathbf{A})^{-1}$ $= \mathbf{V}(\Sigma_{tr}^\dagger)^2\mathbf{V}^T$
\mathbf{r} -estimator	$(\mathbf{I}_m - \mathbf{A}(\mathbf{A}^T\mathbf{A})^{-1}\mathbf{A}^T)\mathbf{b}$	$(\mathbf{I}_m - \mathbf{A}(\mathbf{A}^T\mathbf{A} + \lambda^2\mathbf{I}_n)^{-1}\mathbf{A}^T)\mathbf{b}$	$(\mathbf{I}_m - \mathbf{A}(\mathbf{A}^T\mathbf{A})^{-1}\mathbf{A}^T\mathbf{U}_{tr}\mathbf{U}_{tr}^T)\mathbf{b}$
\mathbf{C}_r	$(\mathbf{I}_m - \mathbf{A}(\mathbf{A}^T\mathbf{A})^{-1}\mathbf{A}^T)$ $= \mathbf{I}_m - \mathbf{U}_1\mathbf{U}_1^T$	$(\mathbf{I}_m - \mathbf{A}(\mathbf{A}^T\mathbf{A} + \lambda^2\mathbf{I}_n)^{-1}\mathbf{A}^T)$ $= \mathbf{I}_m - \mathbf{U}_1\mathbf{F}^2\mathbf{U}_1^T$	$(\mathbf{I}_m - \mathbf{A}(\mathbf{A}^T\mathbf{A})^{-1}\mathbf{A}^T\mathbf{U}_{tr}\mathbf{U}_{tr}^T)$ $= \mathbf{I}_m - \mathbf{U}_{tr}\mathbf{U}_{tr}^T$
Mean	$\mathbf{0}$	$(\mathbf{I}_m - \mathbf{A}(\mathbf{A}^T\mathbf{A} + \lambda^2\mathbf{I}_n)^{-1}\mathbf{A}^T)\mathbf{A}\mathbf{x}^*$	$(\mathbf{I}_m - \mathbf{A}(\mathbf{A}^T\mathbf{A})^{-1}\mathbf{A}^T\mathbf{U}_{tr}\mathbf{U}_{tr}^T)\mathbf{A}\mathbf{x}^*$
Variance	$\mathbf{I}_m - \mathbf{A}(\mathbf{A}^T\mathbf{A})^{-1}\mathbf{A}^T$ $= \mathbf{I}_m - \mathbf{U}_1\mathbf{U}_1^T$	$(\mathbf{I}_m - \mathbf{A}(\mathbf{A}^T\mathbf{A} + \lambda^2\mathbf{I}_n)^{-1}\mathbf{A}^T)^2$ $= \mathbf{I}_m - \mathbf{U}_1\mathbf{F}^2\mathbf{U}_1^T$	$(\mathbf{I}_m - \mathbf{A}(\mathbf{A}^T\mathbf{A})^{-1}\mathbf{A}^T\mathbf{U}_{tr}\mathbf{U}_{tr}^T)$ $= \mathbf{I}_m - \mathbf{U}_{tr}\mathbf{U}_{tr}^T$

Therefore the estimator of $\boldsymbol{\eta} = \mathbf{b} - \mathbf{A}\mathbf{x}^*$ is biased if $\mathbf{F} \neq \mathbf{I}_n$. Similarly, the variance is

$$\mathbf{C}\mathbf{C}^T = \mathbf{U}(\mathbf{I}_m - \mathbf{F})^2\mathbf{U}^T \quad (4.54)$$

$$= \mathbf{U}_2\mathbf{U}_2^T + \mathbf{U}_1(\mathbf{I}_n - \mathbf{F})^2\mathbf{U}_1^T. \quad (4.55)$$

Therefore, in return for a much lower variance for the estimator of \mathbf{x}^* obtained by using regularization, we settle for biased estimators of \mathbf{x}^* and $\boldsymbol{\eta}$ and a somewhat higher variance for the $\boldsymbol{\eta}$ estimator.

We summarize these results in table 1, and in table 2 we write all of the expressions in terms of the data \mathbf{A} and \mathbf{b} , although the expressions in table 1 are both more revealing and more computationally tractable.

5. Choosing a regularization parameter

From (4.28) and (4.34), it can be seen that the residual norm is a monotonically decreasing function of the regularization parameter p for the TSVD method and a monotonically increasing function of the regularization parameter λ for Tikhonov's method. Similarly, the norm of the estimated solution is a monotonically increasing function of p and a monotonically decreasing function of λ . We discuss in this section various alternative methods for choosing the regularization parameter.

5.1. Well-known methods

There are three common methods for choosing the regularization parameter.

The most basic choice of regularization parameter is via *Morozov's discrepancy principle* [22]. This method chooses the regularization parameter so that the norm of the residual is approximately equal to its expected value, given in (2.10).

An alternate strategy for Tikhonov regularization is to choose λ to minimize the *generalized cross-validation* (GCV) function

$$G(\lambda) = \sum_{k=1}^m [b_k - (\mathbf{A}\tilde{\mathbf{x}}_\lambda^{(k)})_k]^2, \quad (5.56)$$

where $\tilde{\mathbf{x}}_\lambda^{(k)}$ is the estimate when the k th measurement b_k is omitted. The basic idea, first introduced by Wahba [34], is to choose λ to make $\mathbf{A}\tilde{\mathbf{x}}_\lambda$ the best overall predictor for missing data values. The formulation for TSVD is similar, but in this case $G(p)$ is a discrete function rather than a continuous one.

A third strategy, often used when the error cannot be assumed to be normally distributed, is based on the *L-curve* first introduced by Hanson and Lawson [19] and further developed by Hansen [9] and [11, chapter 4]. In current use, the L-curve for Tikhonov regularization is a plot of $\log \|\tilde{\mathbf{x}}_\lambda\|$ versus $\log \|\mathbf{b} - \mathbf{A}\tilde{\mathbf{x}}_\lambda\|$ for a range of positive values of λ . The graph on the left of figure 5 is an L-curve plot for the modified Phillips problem. The λ chosen corresponds to the point in the corner where the curvature of the curve is maximized. A method for calculating this point has been described by Hansen and O'Leary [13], and Hansen [10] provides MATLAB software for doing the calculations. The L-curve for TSVD is defined in a similar way, but consists of discrete points for various values of p ; curvature can be defined as the curvature of a smoothing cubic spline.

5.2. Diagnostics based on statistical properties of the residual and its periodogram

Rust [25] (see also [26]) suggested several diagnostics for judging the acceptability of a regularized solution $\tilde{\mathbf{x}}$ with residual $\tilde{\mathbf{r}}$ using properties of the true residual $\boldsymbol{\eta} = \mathbf{r}^*$.

Diagnostic 1. The residual norm-squared should be within two standard deviations of the expected value of $\|\boldsymbol{\eta}\|^2$, which is m ; in other words, $\|\tilde{\mathbf{r}}\|^2 \in [m - 2\sqrt{2m}, m + 2\sqrt{2m}]$. This is a way to quantify the Morozov discrepancy principle.

Diagnostic 2. The elements of $\boldsymbol{\eta}$ are drawn from a $N(0, 1)$ distribution, and a graph of the elements of $\tilde{\mathbf{r}}$ should look like samples from this distribution. (In fact, a histogram of the entries of $\tilde{\mathbf{r}}$ should look like a bell curve.)

Diagnostic 3. We consider the elements of both $\boldsymbol{\eta}$ and $\tilde{\mathbf{r}}$ as time series, with the index i ($i = 1, \dots, m$) taken to be the time variable. Since $\boldsymbol{\eta} \sim N(\mathbf{0}, \mathbf{I}_m)$, the η_i form a white noise series. Therefore the residuals $\tilde{\mathbf{r}}$ for an acceptable estimate should constitute a realization of such a series.

A formal test of diagnostics 2 and 3, used by Rust [25], is based on a plot of the *periodogram* [5, chapter 7], which is an estimate of the power spectrum on the frequency interval $0 \leq f \leq \frac{1}{2T}$, where T is the sample spacing for the time variable. Here, the time variable is the element number i , so $T = 1$. The periodogram is formed by zero-padding the time series to a convenient length N (e.g., a power of 2), taking the discrete Fourier transform of this zero-padded series, and taking the squares of the absolute values of the first half of the coefficients. This gives us the periodogram \mathbf{z} of values corresponding to the frequencies

k/NT , $k = 0, \dots, N/2$. The *cumulative periodogram* \mathbf{c} is the vector of partial sums of the periodogram, normalized by the sum of all of the elements; i.e.,

$$c_k = \frac{\sum_{j=0}^k z_j}{\sum_{j=0}^{N/2} z_j}, \quad (5.57)$$

where \mathbf{z} and \mathbf{c} are vectors with $(N + 1)$ components. The variance in a white noise record is distributed uniformly on the interval $0 \leq f \leq \frac{1}{2T}$.⁵ Of course, no real noise record could have such an even distribution of variance, but on average over many samples, the distribution would approach that uniformity, and in the limit the periodogram would be flat. This means that a plot of the elements of \mathbf{c} versus their frequencies $f_k = k/NT$ would be a straight line passing through the origin with slope $2/NT$. Since the Fourier transform of the residuals is obtained by a linear transformation with an orthogonal matrix, the real and imaginary parts of the transform would remain independently, normally distributed. This implies that the periodogram ordinates would be multiples of independent $\chi^2(2)$ samples and hence that the c_k would be distributed like an ordered sample of size $N/2$ from a uniform $(0,1)$ distribution. A test of the hypothesis that the residuals are white noise can be obtained by constructing the two lines parallel to the one above, passing through the points $(f, c) = (0, \pm\delta)$, where δ is the 5% point for the Kolmogorov–Smirnov statistics for a sample of size $m/2$. These two lines define a 95% confidence band for white noise. More details on this procedure can be found in Fuller [5, pp 363–6].

Therefore, the ideal cumulative periodogram is a straight line between 0 and 1 as the frequency varies between 0 and 0.5, so we expect its length to be close to 1.118 (taking $T = 1$). Quantitative measures for diagnostics 2 and 3 can be based on the length of the cumulative periodogram (estimated as the length of the piecewise linear interpolant to the cumulative periodogram) and on the number of samples outside the 95% confidence band. These measures, along with the residual norm-squared for diagnostic 1, comprise our numerical diagnostics, which we use in conjunction with the plots of the residual vector, its periodogram and its cumulative periodogram.

As an example, we apply these diagnostics to the linear regression estimate $\hat{\mathbf{x}}$ (see (3.13)) for the modified Phillips problem discussed in section 3. Using these diagnostics, we see that this estimate of the solution is not reasonable. We have already observed that $\|\hat{\mathbf{r}}\|^2 = 43.01$, which is outside the $\pm 2\sigma$ confidence interval $[251.02, 348.98]$, so diagnostic 1 is violated. The components of the residual $\hat{\mathbf{r}}$ are plotted against i in the middle left in figure 2, with their histogram displayed in the bottom left. They do not look like samples from a $N(0, 1)$ distribution, so diagnostic 2 is violated. By considering plots of the periodogram and cumulative periodogram on the right in the figure, we see that diagnostic 3 is violated. The middle-right graph in figure 2 is a plot of a periodogram estimate at 4097 equally spaced frequency points on the interval $[0, 0.5]$. For white noise, the variance would be distributed uniformly, but there is much more power on $[0, 0.25]$ than on $[0.25, 0.5]$ so the residuals are probably not white noise. The cumulative periodogram is the solid curve at the bottom right of figure 2. The title gives the length of this curve, which can be compared with the value 1.11803 expected for pure white noise. The dashed lines enclose a 95% confidence band for white noise. We expect that the ordinates for a white noise series should lie outside this band for at most 5% of the frequencies. Since 2803 of the 4096 lie outside the band, it is clear that the least squares residuals are not white noise. Thus, the linear regression estimate fails all three diagnostics and is not an acceptable solution. For comparison, the true solution \mathbf{x}^* with residual $\boldsymbol{\eta}$ passes all three diagnostics, as illustrated in figure 4.

⁵ This interval defines the *Nyquist band*, and the power spectrum for white noise is constant at all frequencies in this band, since the autocorrelation function for white noise is zero for all lags except lag 0.

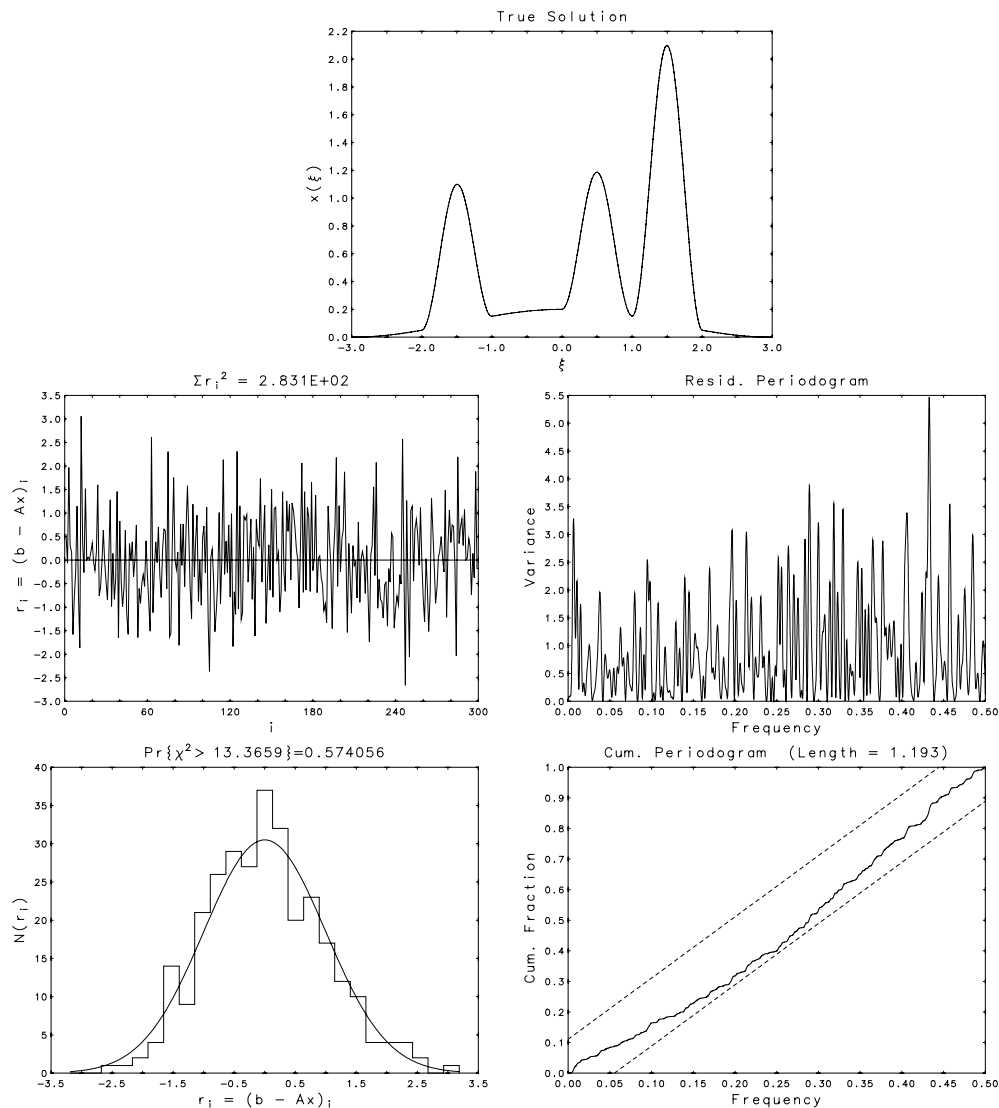


Figure 4. True solution (top) and diagnostics for the modified Phillips problem. On the left are the residual and its histogram. On the right are the residual periodogram and the cumulative periodogram.

In [25], Rust suggested choosing a parameter that passed all three of the diagnostics given above. Later Hansen, Kilmer and Kjeldsen [12] proposed choosing the parameter by either of two methods.

- Choosing the most regularized solution estimate for which the cumulative periodogram lies within the 95% confidence interval.
- Minimizing the sum of the absolute values of the difference between components of \mathbf{c} and the straight line passing through the origin with slope $2/NT$.

The first method tends to undersmooth, since we would expect 5% of the samples to lie outside the confidence interval. For similar reasons, the second method is also too stringent.




Figure 5. L-curve (left) and GCV function (right) for Tikhonov estimates to the solution to the modified Phillips problem. For the L-Curve, the parameter λ (not indicated in the plot) increases from the left to right along the curve, and the ‘corner’ occurs at $\lambda \approx 1.612$. In contrast, the minimizer for the GCV function is $\lambda \approx 13.268$.

Their definition of cumulative periodogram is also somewhat different; they omit the first component of the vector \mathbf{c} and omit zero-padding. Since the first component captures the bias in the estimate, we think it is important to include it. Zero-padding is a standard practice used to give a finer frequency mesh for the periodogram so we prefer to use it. Mead [20, 21] has also recently used distribution properties of the residual to choose regularization parameters.

6. Comparison of parameter choices for the modified Phillips problem

In this section we illustrate the parameter choice methods for Tikhonov regularization and for TSVD using the modified Phillips problem.

6.1. Results for Tikhonov regularization

Consider solving the modified Phillips problem using Tikhonov regularization. For this problem, the ‘optimal’ value from the L-curve in figure 5 is $\lambda = 1.612$. The corresponding estimate and diagnostic plots are given in figure 6, and the numerical diagnostics are given in the first row of table 3. The sum of squared residuals 177.3 is about 5 standard deviations smaller than the expected value 300, and the jaggedness of the plotted solution estimate indicates that a larger λ is needed to provide more smoothing to damp out the higher frequencies so apparent also in the periodogram of the residuals.

The graph on the right-hand side of figure 5 is a plot of the GCV function $G(\lambda)$ versus λ for the modified Phillips problem. The minimum occurs at $\lambda = 13.268$, and the corresponding estimate and residual diagnostics are plotted in figure 7. The numerical diagnostics are tabulated in row 2 of table 3. The fourth column gives the square of the residual norm, which can be compared with $m = 300$. The fifth notes the percentage of cumulative periodogram components within the 95% confidence interval. The sixth column records the length of the graph of the cumulative periodogram, while the next column gives the root mean square error

$$|\Delta x|_{\text{rms}} \equiv \sqrt{\frac{1}{n} \sum_{j=1}^n (\tilde{x}_j - x_j^*)^2}, \quad (6.58)$$

where $\tilde{\mathbf{x}}$ is the computed estimate. (The root-mean-square error is not, of course, computable in practice, since \mathbf{x}^* would be unknown.) Although the estimate is superior to the one for the

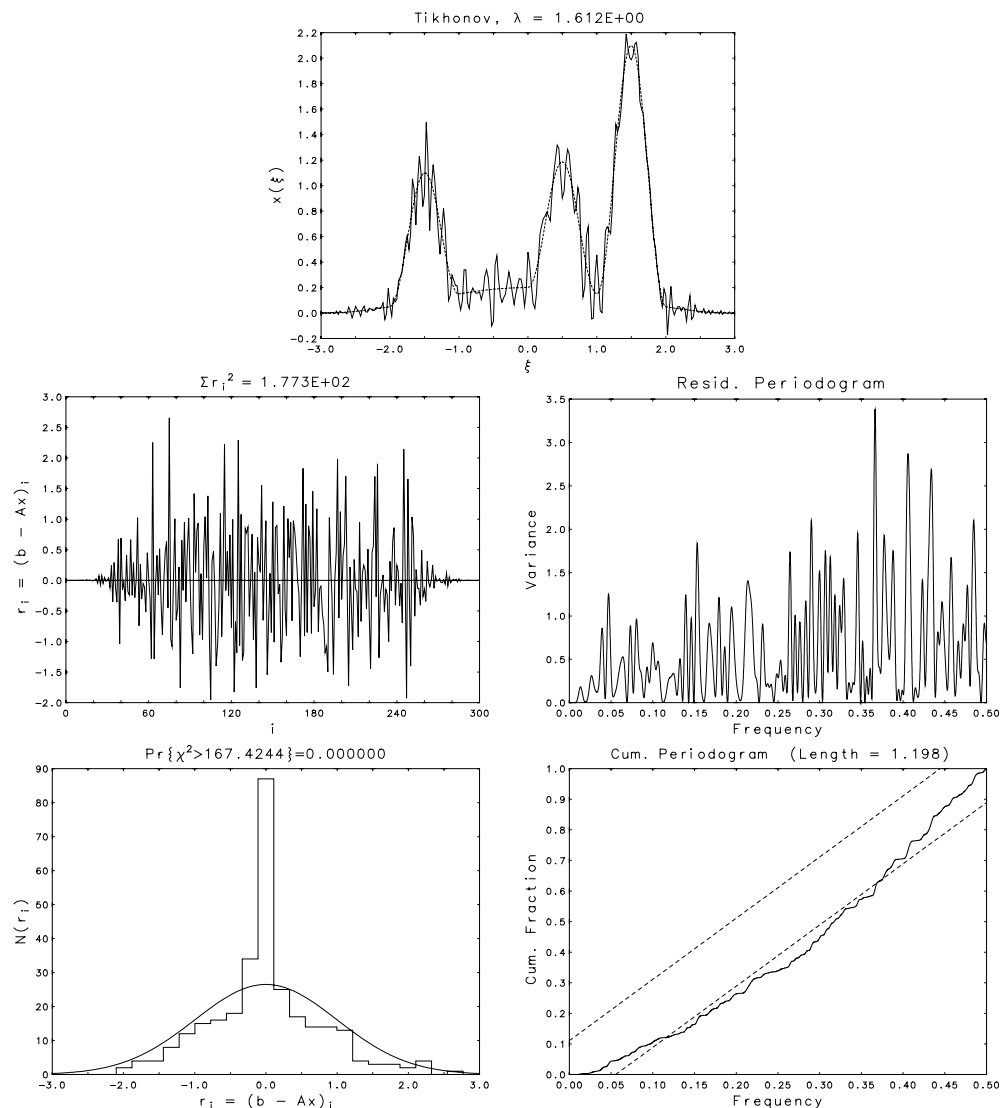


Figure 6. Tikhonov estimate and diagnostics for maximum curvature of the L-curve.

L-curve, it nevertheless does not give acceptable residual diagnostics. The sum of squared residuals is about 3.5 standard deviations smaller than the expected value 300, and 37.7% of the residual cumulative periodogram ordinates lie outside the 95% band for white noise.

Although neither the L-curve λ nor the minimal GCV λ gives an acceptable estimate, it is nevertheless possible to find a good Tikhonov estimate. Rows 3–11 of table 3 give numerical diagnostics for a range of λ s chosen by trial and error to give residuals which have either an acceptable frequency distribution or a sum of squared residuals which differs by less than one standard deviation from the expected value, i.e., to be in the interval $[275.5, 324.5]$. Though all of the values in $19.2 \leq \lambda \leq 27.0$ give acceptable residual frequency distributions, most of them give unacceptably small values for the sum of squared residuals. Only the interval $26.0 \leq \lambda \leq 27.0$ gives estimates with residuals that satisfy all three diagnostics. The estimate

Table 3. Estimate diagnostics for Tikhonov regularization applied to the modified Phillips Problem

	Method	λ	$\sum(\mathbf{b} - \mathbf{A}\tilde{\mathbf{x}})_i^2$	% ($c_k \in \mathcal{B}_{95}$)	Length{c}	$ \Delta\mathbf{x} _{\text{rms}}$	Figure
1	L-curve	1.61	177.3	48.4	1.1982	0.121 31	6
2	GCV	13.27	214.1	62.3	1.2004	0.012 71	7
3	Trial λ	19.20	233.1	95.3	1.1991	0.009 48	
4	Trial λ	20.00	236.8	99.1	1.1992	0.009 25	
5	Trial λ	22.00	247.5	100.0	1.1998	0.008 80	
6	Trial λ	24.00	260.7	100.0	1.2014	0.008 53	
7	Trial λ	26.00	276.9	100.0	1.2040	0.008 41	
8	Trial λ	27.00	286.3	96.9	1.2057	0.008 39	8
9	Trial λ	28.00	296.6	87.5	1.2076	0.008 41	
10	Trial λ	29.00	307.9	83.9	1.2098	0.008 44	
11	Trial λ	30.00	320.2	76.8	1.2123	0.008 50	

for $\lambda = 27.0$, shown in figure 8, is chosen as the ‘optimal’ Tikhonov estimate. It is interesting that even though only 3.1% of the cumulative periodogram ordinates lie outside the 95% white noise band, the cumulative periodogram curve is longer than many of those corresponding to unacceptable estimates. The cause is the low-frequency ‘bulge’ in the distribution which arises from the Tikhonov damping of the middle and high frequencies. Too much damping causes this bulge to burst out of the 95% band and renders the estimate unacceptable.

6.2. Results for TSVD

There are only n possible TSVD estimates, corresponding to the n possible choices for the parameter p in (4.25), (although interpolation can be used to make the solution a continuous piecewise linear function of the parameter p). We show the discrete L-curve for the modified Phillips problem on the left in figure 9 and the discrete GCV function on the right. The corner of the discrete L-curve occurs at $p = 98$ and the minimum of the GCV function occurs at $p = 65$. The numerical diagnostics for these two estimates are given in rows 1 and 3 of table 4. Neither the L-curve nor the GCV estimate gives acceptable residuals. The sum of squared residuals for the L-curve estimate is about 4 standard deviations below the expected value $m = 300$, and for the minimum GCV estimate, about 2.8 standard deviations too small. Only the estimates for $p = 55$ and $p = 54$ give residuals with both a sum of squares in the 1-sigma interval $[275.5, 324.5]$ and an acceptable frequency distribution, with at least 95% of the coefficients within the confidence limits. The $p = 54$ estimate was chosen as optimal, though it is only marginally better than the $p = 55$ estimate. The estimated solution is very much like that obtained by Tikhonov regularization and displayed in figure 8. Note that although the GCV residual is far from appearing to be white noise, for this problem GCV gives the smallest value of $|\Delta\mathbf{x}|_{\text{rms}}$. (We will see evidence in section 8 that this is not typical.) We note that none of the methods found the value $p = 67$ that minimizes the error in \mathbf{x} .

7. Automating the residual periodogram diagnostics

We have seen that the plots of the residual vector, its periodogram, and its cumulative periodogram, in conjunction with the three diagnostics in section 5.2, lead to good choices of a regularization parameter, and we advise their use whenever possible. However, in some circumstances (e.g., real-time applications), it is not feasible to choose the parameter manually, so it is also important to have automatic procedures. There are several ways in which the

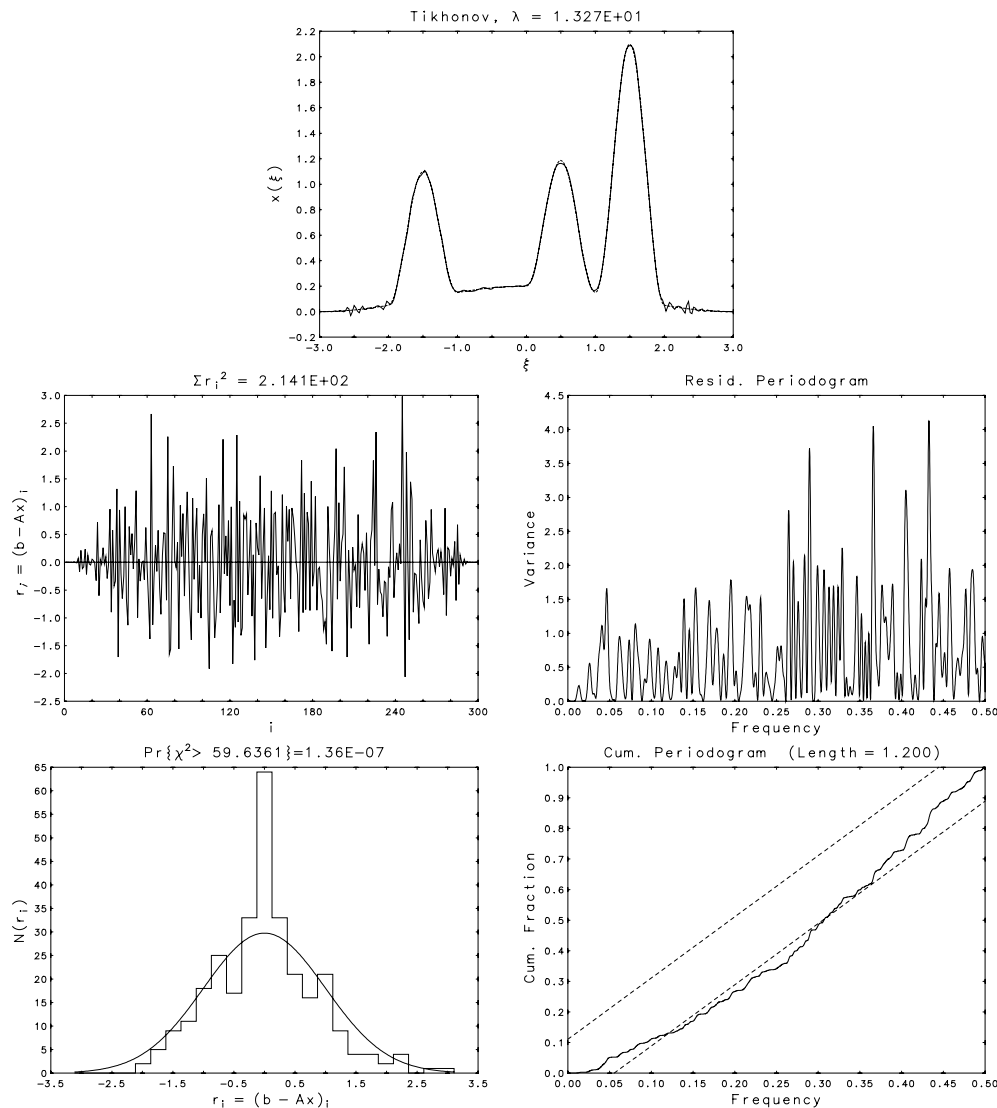


Figure 7. Tikhonov estimate and diagnostics for minimum value of the GCV function.

periodogram can be used: for example, in accord with diagnostics 2 and 3, we can choose the parameter according to any of the following rules, or any combination of them.

- Minimize the length of the cumulative periodogram.
- Minimize the number of points of the cumulative periodogram that lie outside the 95% confidence limits.
- Choose the smallest amount of regularization that places 95% of the points of the cumulative periodogram inside the confidence limits.
- Minimize the distance between the cumulative periodogram and the line through the origin with slope $2/NT$, using either the 1-, 2-, or ∞ -norm to measure distance.
- Choose a value for which the maximum magnitude element in the periodogram is not significantly bigger than expected.

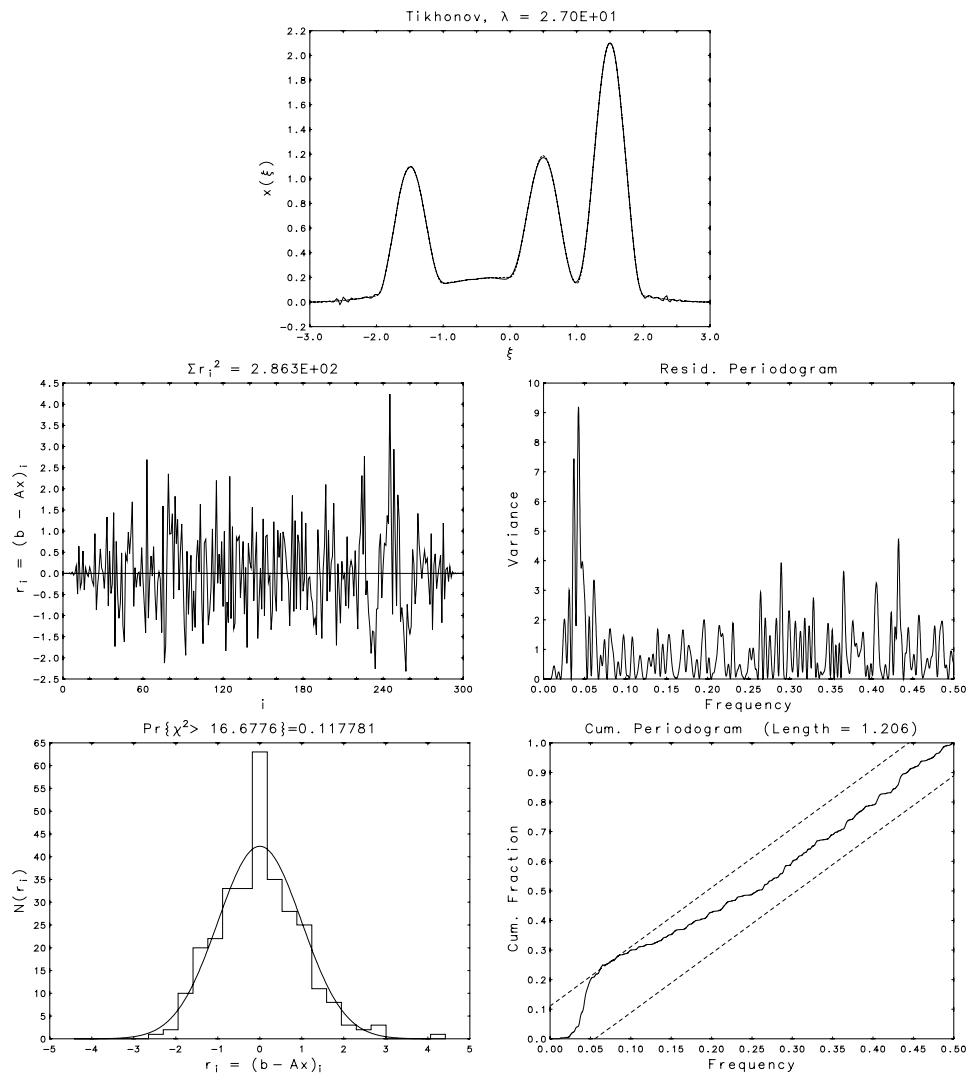


Figure 8. Tikhonov estimate with regularization parameter chosen to optimize the residual diagnostics.

It is also a good idea to restrict our choice to one which puts the norm of the residual within 2 standard deviations of its expected value m (diagnostic 1).

In the following section, we experiment with the last of the options for using the periodogram, by computing the ratio of the largest value to the average value. The distribution of this quantity has been studied by Fisher (see, for example, [5, p 363]), and we chose the least amount of regularization for which the probability of a sample from this distribution being larger than our ratio was at least 5%.

8. Results on some standard test problems

We tested several parameter choice methods on some standard test problems taken from the MATLAB regularization tools [10]: Baart, Foxgood, Heat, Ilaplace, Wing, Phillips, Shaw and




Figure 9. L-curve (left) and GCV function (right) for TSVD estimates for the modified Phillips problem.

Table 4. Diagnostics for TSVD estimates to the solution of the modified Phillips problem.

	Method	p	$\sum (\mathbf{b} - \mathbf{A}\tilde{\mathbf{x}})_i^2$	$\%(c_k \in \mathcal{B}_{95})$	Length $\{\mathbf{c}\}$	$ \Delta x _{\text{rms}}$
1	L-curve	98	201.8	53.5	1.1995	0.035 22
2	Trial p	80	211.7	52.5	1.1964	0.017 51
3	GCV	65	230.8	64.1	1.2032	0.008 71
4	Opt. x	67	229.3	63.2	1.2026	0.008 15
5	Trial p	62	241.5	83.6	1.2013	0.009 43
6	Trial p	61	241.5	82.9	1.2013	0.009 32
7	Trial p	60	244.2	79.5	1.2020	0.008 73
8	Trial p	59	250.6	84.7	1.2020	0.009 61
9	Trial p	58	250.8	82.3	1.2023	0.009 52
10	Trial p	57	259.9	100.0	1.2020	0.010 68
11	Trial p	56	259.9	100.0	1.2021	0.010 68
12	Trial p	55	281.3	100.0	1.2055	0.012 24
13	Trial p	54	281.6	100.0	1.2056	0.012 22
14	Trial p	53	350.9	100.0	1.2091	0.014 02
15	Trial p	52	358.5	100.0	1.2095	0.014 00

Spikes. We generated 100 white noise samples for each of the eight problems with standard deviation equal to 0.001 times the norm of the right-hand side, and used $m = n = 256$.⁶ Table 5 shows the results. We tabulate the number of problems for which the relative error in the regularized solution \mathbf{x}_λ was less than 20%; i.e.,

$$\frac{\|\mathbf{x}_\lambda - \mathbf{x}^*\|}{\|\mathbf{x}^*\|} < .2, \quad (8.59)$$

There were no acceptable solutions from either the Tikhonov or the TSVD family for the problems Wing and Spikes (which both have discontinuous solutions), so we omit those problems from the table. The last column shows that in all the problems we consider, there is a Tikhonov parameter that produces a solution with low relative error and similarly for TSVD. The other table entries show how often each method finds such a parameter. The discrepancy

⁶ Note that the original version of Ilaplace generated a matrix with undefined entries, but this was corrected in an updated version.

Table 5. Number of solutions with relative error less than 20%.

Problem	Method	PFT	Discr.	GCV	L-Curve	HKK	Optimal
Bart	Tikhonov	96	59	67	75	98	100
	TSVD	93	58	70	37	97	100
Foxgood	Tikhonov	98	58	81	76	99	100
	TSVD	96	58	79	35	99	100
Heat	Tikhonov	91	84	86	4	99	100
	TSVD	67	62	86	21	35	100
Ilaplace	Tikhonov	95	68	81	88	99	100
	TSVD	96	68	81	78	99	100
Phillips	Tikhonov	100	91	100	87	100	100
	TSVD	99	89	96	37	97	100
Shaw	Tikhonov	99	59	80	79	99	100
	TSVD	96	60	88	69	98	100
Overall	Tikhonov	579	419	495	409	594	600
		(97%)	(70%)	(83%)	(68%)	(99%)	(100%)
	TSVD	547	395	500	277	525	600
		(91%)	(66%)	(83%)	(46%)	(88%)	(100%)

principle (Discr.), which chooses the parameter for which the residual norm is closest to its expected value, is a rather reliable method, producing a solution with low relative error for Tikhonov on 70% of the examples and for TSVD on 66% of the examples. The GCV works even better, with an acceptable solution in 83% of the examples, despite the fact that the GCV function is notorious for being flat, making it difficult to find a minimizer. The L-curve is slightly less reliable than the discrepancy principle, but still works well in about 68% of the problems for Tikhonov and 46% for TSVD. Often it fails due to a bad estimate of the location of the corner, since it is easily fooled by roughness in the curve. The periodogram with Fisher test (PFT) is quite reliable, finding an acceptable solution for 97% of the problems when using the Tikhonov method and 91% when using TSVD. This is comparable to the results for the method of Hansen, Kilmer and Kjeldsen (HKK), but their method has the statistical flaws of ignoring the first component of the periodogram and of demanding that 100% of the cumulative periodogram values lie within a 95% confidence interval. Based on these and other tests, we conclude that many variants of parameter choice methods based on residual periodograms are quite effective.

9. A real-world example

In the previous section we demonstrated the effectiveness of using the Fisher test on the periodogram in order to automatically determine a regularization parameter. In this section we demonstrate how a spectroscopy problem can be solved using manual application of the diagnostics.

Measurements of nuclear radiation spectra are an important source of ill-posed problems. Consider the energy spectrum of neutrons produced by the reaction $T(d, n)^4\text{He}$, i.e. tritium nuclei bombarded with deuterons to produce helium nuclei and the neutrons being measured. If the bombarding particles are monoenergetic then so are the neutrons produced. Although the neutrons are monoenergetic, the measuring instrument both smears and distorts the expected spectrum. The instrument in question, an NE-213 spectrometer, has been described by Verbinski *et al* [33] and Burrus and Verbinski [3]. Its response functions are plotted in figure 10, where incident energy corresponds to the variable ξ and pulse height to the variable




Figure 10. Instrument response functions for the Burrus neutron spectrum problem. The quantity $\log_{10}[1 + K_i(E)]$ is plotted in order to more clearly show the structure for higher energies.

t_i in (1.1). For each neutron, the detector increments a counter in one of $m = 113$ pulse-height bins designed to cover and subdivide the energy range of all possible incident neutrons. The integral equations modeling the process are

$$c_i = \int_0^{E_{up}} K_i(E) N(E) dE + \varepsilon_i, \quad i = 1, 2, \dots, 113, \quad (9.60)$$

where c_i is the number of pulses counted in the i th bin, $N(E)$ is the unknown number of neutrons at energy E and the $K_i(E)$ are the instrument response functions for the detector⁷. For a given E , the quantity $K_i(E) dE$ is the probability that a neutron with energy in the range $E \pm \frac{1}{2} dE$ would increment the count in the i th pulse-height bin. Ideally, higher energy neutrons would be counted in higher bins, but figure 10 shows that this does not always happen. Ideally, the $K_i(E)$ should be a series of narrow peaks distributed along a diagonal line running from the upper left to the lower right of the energy, pulse-height domain. The figure does exhibit a ridge along that direction, but it attenuates to a barely discernible ripple for higher energies. Even worse, the long energy tails for the lower pulse-height bins make it more likely that a higher energy neutron will produce an increment in a lower rather than a higher bin.

Figure 11 shows a measured spectrum with count rates rather than counts plotted against pulse height. Counts are accumulated for a time sufficient to allow good estimates for the uncertainties and then are divided by the elapsed time to get rates. This division replaces the c_i with count rates y_i and the number $N(E)$ of neutrons with a neutron flux $x(E)$. As a result, (9.60) becomes

$$y_i = \int_0^{E_{up}} K_i(E) x(E) dE + \epsilon_i, \quad i = 1, 2, \dots, 113, \quad (9.61)$$

where the ϵ_i are noise rates derived from dividing estimates of the measurement errors ε_i by time. These errors are modeled as samples from a Poisson distribution, with parameter equal

⁷ The data for this problem are posted at www.cs.umd.edu/users/oleary/bwr.




Figure 11. Measured pulse-height spectrum for the Burrus neutron spectrum problem. The two curves are ± 1 -sigma bounds for the measurements.




Figure 12. Singular values and first n elements of $|U^T b|$ for the Burrus neutron spectrum problem. The vertical line marks the first deleted singular value (σ_{41}) for the optimal TSVD estimate.

Table 6. Estimate diagnostics for the Burrus neutron spectrum problem.

Method	Parameter	$\sum (\mathbf{b} - \mathbf{A}\tilde{\mathbf{x}})_i^2$	$\%(c_k \in \mathcal{B}_{95})$	Length{c}	Figure
1 Least squares	$\lambda = 0$	38.0	37.4	1.255	
2 Tikhonov regularization, L-curve	$\lambda = 61.45$	50.0	51.4	1.250	
3 Tikhonov regularization, GCV	$\lambda = 100.1$	57.6	55.4	1.222	
4 Tikhonov regularization, optimal trial λ	$\lambda = 233.0$	113.4	100.0	1.174	13
5 TSVD, L-curve	$p = 61$	50.0	50.8	1.253	
6 TSVD, GCV	$p = 49$	61.1	51.7	1.256	
7 TSVD, optimal trial λ	$p = 40$	115.2	98.8	1.224	14

to the true number of counts. Running the experiment sufficiently long makes the counts large enough (e.g., at least 50) to approximate the Poisson distribution by a normal distribution with

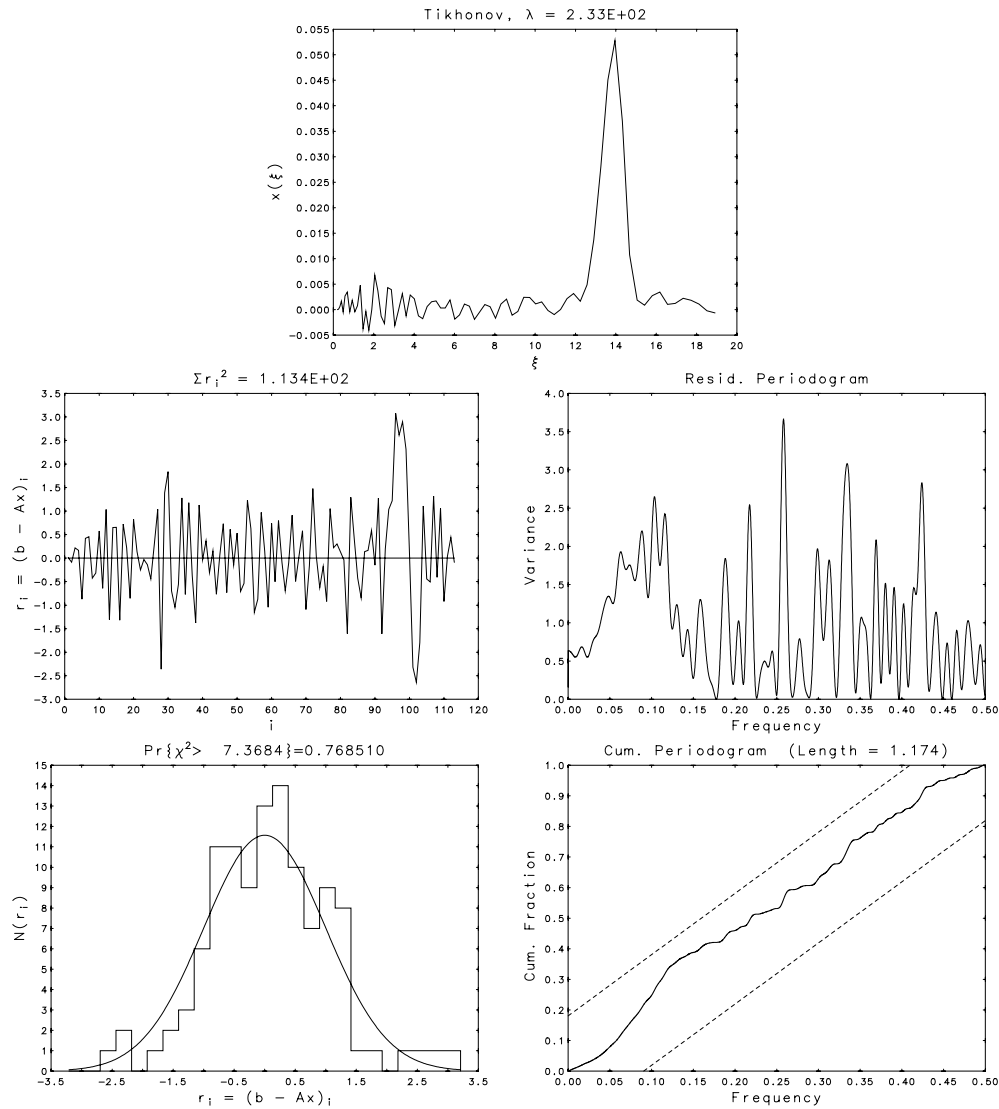


Figure 13. Optimal Tikhonov solution estimate and diagnostic plots for the Burrus neutron spectrum problem.

mean and variance equal to this number [5]. Therefore the i th component of \mathbf{y} has a Poisson distribution with parameter $(\mathbf{K}\mathbf{x})_i$, and we approximate this distribution by $N((\mathbf{K}\mathbf{x})_i, (\mathbf{K}\mathbf{x})_i)$. Since $\mathbf{K}\mathbf{x}$ is unknown, we approximate the parameters by y_i .

Assuming statistical independence and dividing the estimated standard deviations by the time of accumulation gives a diagonal variance matrix \mathbf{S}^2 for the errors.

Equations (9.61) were discretized using simple rectangular quadrature with an $n = 77$ point energy mesh, $0.2\text{MeV} = E_1 < E_2 < \dots < E_{77} = 18.91\text{ MeV}$, to give the linear regression model (1.2). Scaling with \mathbf{S}^{-1} gave the normalized model (2.7). The two curves plotted in figure 11 indicate good statistics for the measurements, but there is no sign of the peak expected for monoenergetic neutrons. The sharp rise in the lower pulse-height bins

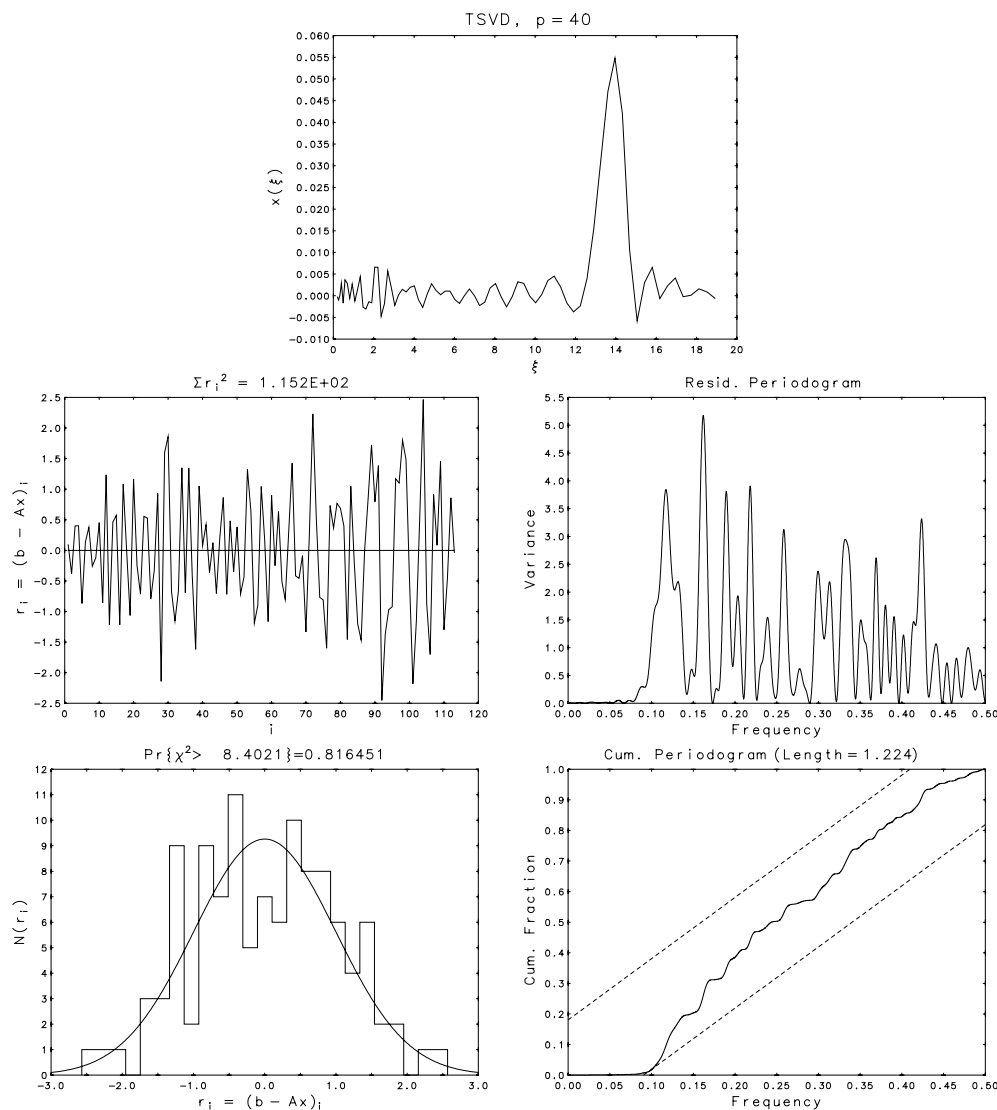


Figure 14. Optimal TSVD solution estimate and diagnostic plots for the Burrus neutron spectrum problem.

contains most of the counts that should have gone into that missing peak. The least squares estimate is overwhelmed by the amplified noise and gave $r_{\min}^2 = \|\mathbf{b} - \mathbf{A}\hat{\mathbf{x}}\|^2 = 38.05$, which is almost 5 standard deviations smaller than the expected value 113. The singular values and first n elements of $|\mathbf{U}^T \mathbf{b}|$ are shown in figure 12. The condition number is 2.3503×10^5 , so the matrix is clearly not rank-deficient. The distribution of the $|\mathbf{U}^T \mathbf{b}|_i$ exhibits a dichotomy at $i = 47$ where it first drops below the 1-sigma level for the error.

Table 6 gives diagnostics for several estimates of the energy spectrum. The quantities tabulated are the same as in the previous tables except that here $|\Delta x|_{\text{rms}}$ is not known. The first row corresponds to the least squares estimate.

Rows 2–4 of table 6 give diagnostics for three Tikhonov estimates. The L-curve method is used in row 2. The sum of squared residuals is about 4.2 standard deviations smaller than

the expected value, and an excess of higher frequencies rendered the frequency distribution of the residuals unacceptable. In row 3, the smoothing constant is chosen to minimize the GCV function. The sum of squared residuals is about 3.7 standard deviations smaller than the expected value, and the frequency distribution of the residuals is unacceptable. For the estimate in row 4, the smoothing constant is chosen by trial and error to give a sum of squared residuals close to the expected value 113. The estimated solution and diagnostic plots are given in figure 13.

Rows 5–7 of table 6 give diagnostics for three estimates obtained from the TSVD method. For rows 5 and 6, the truncation parameters are chosen by the L-curve and minimum GCV methods, respectively. The estimate and diagnostics are very similar to those for the corresponding Tikhonov estimates (rows 1 and 2): the estimated solution is reasonable but the residual is unacceptable. Row 7 gives the diagnostics for the optimal TSVD truncation. The solution estimate and diagnostic plots are given in figure 14.

10. Discussion and conclusions

For ill-posed problems with errors dominated by the measurement uncertainties in \mathbf{y} , the instability in the solution estimate is attributable to components of the measurements which are overwhelmed by those uncertainties. In general, these components do not correspond exactly to specific elements of \mathbf{y} . When possible, we advocate scaling the problem by \mathbf{S}^{-1} to transform it to one in which the errors all have unit variance. Regularization methods can then be used to reduce the influence of the error. Ideally, the results of candidate parameter choices should be evaluated by plotting the resulting residual along with its periodogram and its cumulative periodogram and examining the numerical diagnostics. Sometimes an automated choice is needed, and in such cases we advise using diagnostics based on the norm of the residual and the periodogram of the residual to choose or validate the choice of a regularization parameter.

To apply either the manual or the automatic methods for parameter choice, one need only be able to compute the Fourier coefficients of the (zero-padded) residual for various values of the regularization parameter. Thus all of these numerical diagnostics are inexpensive to compute, so they can be used for two- and three-dimensional problems, even if regularization methods other than the SVD-based ones discussed here are used. They can be used, for example, in conjunction with iterative regularization methods [18] or even with nonlinear models.

Of course, there is no guarantee that the estimate resulting from our parameter choice method will be close to the true solution, since (quite likely) the assumption that the errors η are $N(\mathbf{0}, \mathbf{I}_m)$ may only be an approximation. Even if the assumption is good, noise may be large enough to overwhelm important components of the signal. To recover such components it is necessary to repeat the measurements multiple times or with more accuracy.

Acknowledgments

BWR would like to thank Drs W R Burrus and R E Funderlic for the initial inspiration. We thank Drs A S Carasso, K J Coakley, D E Gilsinn, K A Remington, and the anonymous referees for their reviews and suggestions and Dr M G Vangel for suggestions for improving the statistical aspects of this manuscript. We are grateful to Drs M E Kilmer and P C Hansen for sharing their software. The work of the second author was supported in part by the National Science Foundation under grant CCF 0514213. The views expressed here are those

of the authors alone, not necessarily those of NIST or NSF. Certain commercial products are identified in order to specify adequately experimental procedures. In no case does such identification imply recommendation or endorsement by NIST, nor does it imply that the items identified are necessarily the best available for the purpose.

Appendix A. A variant of the Phillips problem

A useful test problem which shares many of the characteristics of real instrument correction problems is obtained by discretizing a variant of the well-known [23] Phillips equation. This modified Phillips problem can be written as

$$y(t) = \int_{-3}^3 K(t, \xi) x(\xi) d\xi, \quad -6 \leq t \leq 6, \quad (\text{A.1})$$

with the kernel given by

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

and the exact solution by

$$x(\xi) = \beta(\xi) + \sum_{k=1}^3 c_k(\xi), \quad (\text{A.3})$$

where

$$\beta(\xi) = \begin{cases} \alpha_0 \left[1 + \cos \left(\frac{\pi}{3} \xi \right) \right], & |\xi| \leq 3, \\ 0, & |\xi| > 3, \end{cases} \quad (\text{A.4})$$

and

$$c_k(\xi) = \begin{cases} \alpha_k \{ 1 + \cos[2\pi(\xi - \psi_k)] \}, & |\xi - \psi_k| \leq \frac{1}{2}, \\ 0, & \text{otherwise,} \end{cases} \quad (\text{A.5})$$

with amplitude constants α_k and centering constants ψ_k chosen to be

$$\begin{aligned} \alpha_0 &= 0.1, & \alpha_1 &= 0.5, & \alpha_2 &= 0.5, & \alpha_3 &= 1.0, \\ \psi_1 &= -1.5, & \psi_2 &= 0.5, & \psi_3 &= 1.5. \end{aligned} \quad (\text{A.6})$$

The kernel differs from the Phillips original only in the inclusion of the normalizing factor $\frac{1}{6}$ which is added to assure that for any ξ ,

$$\int_{-3+\xi}^{3+\xi} K(t, \xi) dt = 1. \quad (\text{A.7})$$

For a measuring instrument this condition assures that conservation laws are not violated. Plots of $K(t, \xi)$ for 5 representative values of ξ are given on the left in figure 1.

The exact solution to the original Phillips problem appears in a scaled down form as the $\beta(\xi)$ term in the solution to the new problem. The scaling constant α_0 is chosen to reduce the original Phillips solution to the role of a background function on which to superimpose the three discrete peaks represented by the $c_k(\xi)$ terms. The three points $\xi = \psi_k$ are the centers of these peaks and the constants $2\alpha_k$ are their heights above the background. The new solution function is plotted as a dashed line on the right in figure 1.

These changes in the Phillips problem are designed to make it more challenging and more reminiscent of real-world instrument correction problems. Unfortunately, they also make the representation of the function $y(t)$ more complicated. Substituting (A.3) into (A.1) gives

$$y(t) = \int_{-3}^3 K(t, \xi) \beta(\xi) d\xi + \sum_{k=1}^3 \int_{-3}^3 K(t, \xi) c_k(\xi) d\xi, \quad -6 \leq t \leq 6, \quad (\text{A.8})$$

but care must be exercised in evaluating these integrals because $K(t, \xi) = 0$ on half of the rectangular domain $\{(t, \xi) | -6 \leq t \leq 6, -3 \leq \xi \leq 3\}$ and each of the $c_k(\xi)$ is zero everywhere except on the interval $\psi_k - \frac{1}{2} \leq \xi \leq \psi_k + \frac{1}{2}$. The last equation can also be written as

$$y(t) = B(t) + \sum_{k=1}^3 C_k(t), \quad (\text{A.9})$$

where

$$B(t) \equiv \int_{-3}^3 K(t, \xi) \beta(\xi) d\xi, \quad (\text{A.10})$$

and

$$C_k(t) \equiv \int_{\psi_k - \frac{1}{2}}^{\psi_k + \frac{1}{2}} K(t, \xi) c_k(\xi) d\xi, \quad k = 1, 2, 3. \quad (\text{A.11})$$

Evaluating the integral for $B(t)$ gives

$$B(t) = \frac{1}{6} \alpha_0 \left\{ (6 - |t|) \left[1 + \frac{1}{2} \cos\left(\frac{\pi}{3} t\right) \right] + \frac{9}{2\pi} \sin\left(\frac{\pi}{3} |t|\right) \right\}, \quad -6 \leq t \leq 6, \quad (\text{A.12})$$

and the integrals for $C_k(t)$ can be written as

$$C_k(t) = \frac{1}{6} \alpha_k L_k(t), \quad k = 1, 2, 3, \quad (\text{A.13})$$

where

$$L_1(t) = \begin{cases} 0, & -6 \leq t \leq -5, \\ t + 5 + \frac{1}{2\pi} \sin[\pi(2t + 9)] + \frac{3}{\pi} \sin\left[\frac{\pi}{3}(t + 2)\right] \\ \quad + \frac{3}{10\pi} \left\{ \sin[\pi(2t + 8)] - \sin\left[\frac{\pi}{3}(t - 1)\right] \right\} \\ \quad + \frac{3}{14\pi} \left\{ \sin[\pi(2t + 10)] + \sin\left[\frac{\pi}{3}(t + 5)\right] \right\}, & -5 \leq t \leq -4, \\ 1 + \frac{3}{\pi} \left\{ -\sin\left[\frac{\pi}{3}(1 + t)\right] + \sin\left[\frac{\pi}{3}(2 + t)\right] \right\} \\ \quad + \frac{3}{10\pi} \left\{ \sin\left[\frac{\pi}{3}(4 + t)\right] + \sin\left[\frac{\pi}{3}(1 - t)\right] \right\} \\ \quad + \frac{3}{14\pi} \left\{ \sin\left[\frac{\pi}{3}(2 - t)\right] + \sin\left[\frac{\pi}{3}(5 + t)\right] \right\}, & -4 \leq t \leq 1, \\ 2 - t + \frac{1}{2\pi} \sin[\pi(3 - 2t)] - \frac{3}{\pi} \sin\left[\frac{\pi}{3}(t + 1)\right] \\ \quad + \frac{3}{10\pi} \left\{ \sin\left[\frac{\pi}{3}(t + 4)\right] + \sin[\pi(2 - 2t)] \right\} \\ \quad + \frac{3}{14\pi} \left\{ \sin\left[\frac{\pi}{3}(2 - t)\right] + \sin[\pi(4 - 2t)] \right\}, & 1 \leq t \leq 2, \\ 0, & 2 \leq t \leq 6, \end{cases} \quad (\text{A.14})$$

$$L_2(t) = \begin{cases} 0, & -6 \leq t \leq -3, \\ t + 3 + \frac{1}{2\pi} \sin[\pi(2t+5)] + \frac{3}{\pi} \sin\left[\frac{\pi}{3}t\right] \\ \quad + \frac{3}{10\pi} \left\{ \sin[\pi(2t+4)] - \sin\left[\frac{\pi}{3}(t-3)\right] \right\} \\ \quad + \frac{3}{14\pi} \left\{ \sin[\pi(2t+6)] + \sin\left[\frac{\pi}{3}(t+3)\right] \right\}, & -3 \leq t \leq -2, \\ 1 + \frac{3}{\pi} \left\{ \sin\left[\frac{\pi}{3}(1-t)\right] + \sin\left[\frac{\pi}{3}t\right] \right\} \\ \quad + \frac{3}{10\pi} \left\{ \sin\left[\frac{\pi}{3}(2+t)\right] + \sin\left[\frac{\pi}{3}(3-t)\right] \right\} \\ \quad + \frac{3}{14\pi} \left\{ \sin\left[\frac{\pi}{3}(4-t)\right] + \sin\left[\frac{\pi}{3}(3+t)\right] \right\}, & -2 \leq t \leq 3, \\ 4 - t + \frac{1}{2\pi} \sin[\pi(7-2t)] + \frac{3}{\pi} \sin\left[\frac{\pi}{3}(1-t)\right] \\ \quad + \frac{3}{10\pi} \left\{ \sin\left[\frac{\pi}{3}(t+2)\right] - \sin[\pi(2t-6)] \right\} \\ \quad + \frac{3}{14\pi} \left\{ \sin\left[\frac{\pi}{3}(4-t)\right] - \sin[\pi(2t-8)] \right\}, & 3 \leq t \leq 4, \\ 0, & 4 \leq t \leq 6, \end{cases} \quad (\text{A.15})$$

and

$$L_3(t) = \begin{cases} 0, & -6 \leq t \leq -2, \\ t + 2 + \frac{1}{2\pi} \sin[\pi(2t+3)] + \frac{3}{\pi} \sin\left[\frac{\pi}{3}(t-1)\right] \\ \quad + \frac{3}{10\pi} \left\{ \sin[\pi(2t+2)] - \sin\left[\frac{\pi}{3}(t-4)\right] \right\} \\ \quad + \frac{3}{14\pi} \left\{ \sin[\pi(2t+4)] + \sin\left[\frac{\pi}{3}(t+2)\right] \right\}, & -2 \leq t \leq -1, \\ 1 + \frac{3}{\pi} \left\{ \sin\left[\frac{\pi}{3}(2-t)\right] - \sin\left[\frac{\pi}{3}(1-t)\right] \right\} \\ \quad + \frac{3}{10\pi} \left\{ \sin\left[\frac{\pi}{3}(1+t)\right] - \sin\left[\frac{\pi}{3}(t-4)\right] \right\} \\ \quad + \frac{3}{14\pi} \left\{ \sin\left[\frac{\pi}{3}(5-t)\right] + \sin\left[\frac{\pi}{3}(2+t)\right] \right\}, & -1 \leq t \leq 4, \\ 5 - t - \frac{1}{2\pi} \sin[\pi(2t-9)] + \frac{3}{\pi} \sin\left[\frac{\pi}{3}(2-t)\right] \\ \quad + \frac{3}{10\pi} \left\{ \sin\left[\frac{\pi}{3}(t+1)\right] - \sin[\pi(2t-8)] \right\} \\ \quad + \frac{3}{14\pi} \left\{ -\sin\left[\frac{\pi}{3}(t-5)\right] - \sin[\pi(2t-10)] \right\}, & 4 \leq t \leq 5, \\ 0, & 5 \leq t \leq 6, \end{cases} \quad (\text{A.16})$$

The function $y(t)$ is plotted as the solid curve on the right in figure 1. All of the details of the three peaks are so smeared together by the convolution with the kernel function that there is no hint of any structure in the underlying $x(\xi)$.

The modified Phillips problem is discretized by choosing $m = 300$ equally spaced mesh points on the interval $-5.9625 \leq t \leq 5.9625$ to give

$$y(t_i) = \int_{-3}^3 K(t_i, \xi) x(\xi) d\xi, \quad i = 1, 2, \dots, 300, \quad (\text{A.17})$$

and by replacing each of the integrals by an $n = 241$ point trapezoidal quadrature sum, i.e.,

$$\int_{-3}^3 K(t_i, \xi) x(\xi) d\xi \approx \sum_{j=1}^{241} \omega_j K(t_i, \xi_j) x(\xi_j), \quad i = 1, \dots, 300, \quad (\text{A.18})$$

where the ξ_j are 241 equally spaced mesh points on the interval $-3.0 \leq \xi \leq 3.0$, and

$$(\omega_1, \omega_2, \omega_3, \dots, \omega_{240}, \omega_{241}) = \frac{1}{2} \frac{6}{n-1} \times (1, 2, 2, \dots, 2, 1). \quad (\text{A.19})$$

Defining the n -vector \mathbf{x}^* and the m -vector \mathbf{y}^* by

$$\begin{aligned} x_j^* &= x(\xi_j), & j &= 1, 2, \dots, 241, \\ y_i^* &= y(t_i), & i &= 1, 2, \dots, 300, \end{aligned} \quad (\text{A.20})$$

and the $m \times n$ matrix \mathbf{K} by

$$K_{i,j} = \omega_j K(t_i, \xi_j), \quad \begin{aligned} i &= 1, 2, \dots, 300, \\ j &= 1, 2, \dots, 241, \end{aligned} \quad (\text{A.21})$$

gives

$$\mathbf{y}^* = \mathbf{K}\mathbf{x}^* + \boldsymbol{\delta}, \quad (\text{A.22})$$

where $\boldsymbol{\delta}$ is an m -vector of quadrature errors. A crucial assumption in replacing the integrals with quadrature sums is that the value of n is chosen large enough so that the δ_i are small relative to the random measuring errors ϵ_i . To assure that this assumption is satisfied for the test problem, the elements of the vector \mathbf{y}^* were computed from the matrix-vector product

$$\mathbf{y}^* = \mathbf{K}\mathbf{x}^* \quad (\text{A.23})$$

rather than from (A.9)–(A.16). More precisely, the matrix elements $K_{i,j}$ were computed from (A.2), (A.19), (A.21), the vector elements x_j^* were computed from (A.3)–(A.6), and the vector \mathbf{y}^* is then computed from (A.23). The ‘measured’ vector \mathbf{y} was then obtained by adding random perturbations to the elements of this \mathbf{y}^* . Each perturbation was chosen independently from a normal distribution with mean zero and standard deviation $s_i = (10^{-5})\sqrt{y_i^*}$, so the variance matrix is

$$\mathbf{S}^2 = \text{diag}(s_1^2, s_2^2, \dots, s_{300}^2), \quad s_i = (10^{-5})\sqrt{y_i^*}. \quad (\text{A.24})$$

References

- [1] Anderson T W 1958 *An Introduction to Multivariate Statistical Analysis* (New York: Wiley)
- [2] Bertero M and Boccacci P 1998 *Introduction to Inverse Problems in Imaging* (Bristol: Institute of Physics Publishing)
- [3] Burrus W and Verbinski V 1969 Fast neutron spectroscopy with thick organic scintillators *Nucl. Instrum. Methods* **67** 181–96
- [4] Engl H, Hanke M and Neubauer A 1996 *Regularization of Inverse Problems* (Dordrecht: Kluwer)
- [5] Fuller W A 1996 *Introduction to Statistical Time Series* (New York: Wiley)
- [6] Golub G and Kahan W 1965 Calculating the singular values and pseudo-inverse of a matrix *J. Soc. Ind. Appl. Math.: Ser. B Numer. Anal.* **2** 205–24
- [7] Golub G H, Klema V and Stewart G W 1976 Rank degeneracy and least squares problems, *Tech. Rep.* STAN-CS-76-559, August 1976, Computer Science Department, Stanford University, CA
- [8] Golub G H and Van Loan C F 1996 *Matrix Computations* 3rd edn (Baltimore: John Hopkins University Press)
- [9] Hansen P C 1992 Analysis of discrete ill-posed problems by means of the L-curve *SIAM Rev.* **34** 561–80
- [10] Hansen P C 1994 Regularization tools: a MATLAB package for analysis and solution of discrete ill-posed problems *Numer. Algorithms* **6** 1–35
- [11] Hansen P C 1998 *Rank-Deficient and Discrete Ill-Posed Problems Numerical Aspects of Linear Inversion* (Philadelphia, PA: SIAM)
- [12] Hansen P C, Kilmer M E and Kjeldsen R H 2006 Exploiting residual information in the parameter choice for discrete ill-posed problems *Biomed. Instrum. Technol.* **46** 41–59
- [13] Hansen P C and O'Leary D P 1993 The use of the L-curve in the regularization of discrete ill-posed problems *SIAM J. Sci. Comput.* **14** 1487–503
- [14] Hanson R J 1971 A numerical method for solving Fredholm integral equations of the first kind using singular values *SIAM J. Numer. Anal.* **8** 616–22
- [15] Hoerl A E and Kennard R W 1970 Ridge regression: Applications to nonorthogonal problems *Technometrics* **12** 69–82
- [16] Hoerl A E and Kennard R W 1970 Ridge regression: biased estimation for nonorthogonal problems *Technometrics* **12** 55–67
- [17] Hogg R V and Craig A T 1965 *Introduction to Mathematical Statistics* 2nd edn (New York: Macmillan)
- [18] Kilmer M E and O'Leary D P 2001 Choosing regularization parameters in iterative methods for ill-posed problems *SIAM J. Matrix Anal. Appl.* **22** 1204–21

- [19] Lawson C L and Hanson R J 1974 *Solving Least Squares Problems* (Englewood Cliffs, NJ: Prentice-Hall)
- Lawson C L and Hanson R J 1995 *Solving Least Squares Problems* (Philadelphia, PA: SIAM) (reprinted)
- [20] Mead J and Renaut R A 2007 The χ^2 -curve method of parameter estimation for generalized Tikhonov regularization *Tech. Rep.*
- [21] Mead J L 2007 Parameter estimation: A new approach to weighting a priori information *J. Inv. Ill-Posed Problems* **15** 1–21
- [22] Morozov V A 1966 On the solution of functional equations by the method of regularization *Sov. Math.—Dokl.* **7** 414–7
- [23] Phillips D L 1962 A technique for the numerical solution of certain integral equations of the first kind *J. Assoc. Comput. Mach.* **9** 84–97
- [24] Rust B and Burrus W 1972 *Mathematical Programming and the Numerical Solution of Linear Equations* (New York: Elsevier)
- [25] Rust B W 2000 Parameter selection for constrained solutions to ill-posed problems *Comput. Sci. Stat.* **32** 333–47
- [26] Rust B W 1998 Truncating the singular value decomposition for ill-posed problems *Tech. Rep. NISTIR 6131* National Institute of Standards and Technology, US Department of Commerce, Gaithersburg, MD
- [27] Smithies F 1958 *Integral Equations (Cambridge Tracts in Mathematics and Mathematical Physics vol. 49)* (New York: Cambridge University Press)
- [28] Tarantola A 1987 *Inverse Problem Theory* (Amsterdam: Elsevier)
- [29] Tikhonov A 1995 *Numerical Methods for the Solution of Ill-Posed Problems* (Berlin: Springer)
- [30] Tikhonov A N 1963 Solution of incorrectly formulated problems and the regularization method *Sov. Math.—Dokl.* **4** 501–4
- [31] Twomey S 1963 On the numerical solution of Fredholm integral equations of the first kind by inversion of the linear system produced by quadrature *J. Assoc. Comput. Mach.* **10** 97–101
- [32] Twomey S 1977 *Introduction to the Mathematics of Inversion in Remote Sensing and Indirect Measurements* (Amsterdam: Elsevier)
- [33] Verbinski V, Burrus W, Love T, Zobel W, Hill N and Textor R 1968 Calibration of an organic scintillator for neutron spectrometry *Nucl. Instrum. Methods* **65** 8–25
- [34] Wahba G 1977 Practical approximate solutions to linear operator equations when the data are noisy *SIAM J. Numer. Anal.* **14** 651–67
- [35] Wing G M 1991 *A Primer on Integral Equations of the First Kind: The problem of deconvolution and unfolding*, With the assistance of J D Zahrt (Philadelphia, PA: SIAM)