

REGULARIZATION AND THE GENERAL GAUSS-MARKOV LINEAR MODEL

HONGYUAN ZHA AND PER CHRISTIAN HANSEN

ABSTRACT. If the coefficient matrix in the general Gauss-Markov linear model is ill-conditioned, then the solution is very sensitive to perturbations. For such problems, we propose to add Tikhonov regularization to the model, and we show that this actually stabilizes the solution and decreases its variance. We also give a numerically stable algorithm for computing the regularized solution efficiently.

1. INTRODUCTION

This paper is concerned with the *general Gauss-Markov linear model* $A\mathbf{x} + \boldsymbol{\varepsilon} = \mathbf{b}$, where $A \in \mathbb{R}^{m \times n}$ ($m \geq n$) and $\mathbf{b} \in \mathbb{R}^m$ are known, $\mathbf{x} \in \mathbb{R}^n$ is an unknown vector to be estimated, and $\boldsymbol{\varepsilon} \in \mathbb{R}^m$ is a random vector with zero mean and variance-covariance matrix $V(\boldsymbol{\varepsilon}) = s^2 BB^T$ with $B \in \mathbb{R}^{m \times p}$ ($m \geq p$). The best linear unbiased estimator of \mathbf{x} in this model is the solution to the following constrained least squares problem:

$$(1) \quad \min \|\mathbf{u}\|_2 \quad \text{subject to} \quad A\mathbf{x} + B\mathbf{u} = \mathbf{b}.$$

Here, we have introduced the vector $\mathbf{u} \in \mathbb{R}^p$ such that $\boldsymbol{\varepsilon} = B\mathbf{u}$, where \mathbf{u} has variance-covariance matrix $V(\mathbf{u}) = s^2 I_p$ and where I_p is the identity matrix. The model (1) was introduced by Paige [10], and computational algorithms can be found in [7, 9, 11]. A more detailed analysis of (1) in terms of the generalized SVD is also given by Paige [12], while Björck [1, §23] extended this analysis to the case when both A and B may be rank-deficient. However, the case when problem (1) is ill-conditioned, for example if A or B is ill-conditioned, has not been given much attention and, according to Paige [10], needs further work. The present paper is a step in this direction.

Received July 12, 1989; revised September 21, 1989.

1980 *Mathematics Subject Classification* (1985 Revision). Primary 65F20, 65F30.

Key words and phrases. Estimation in the general Gauss-Markov linear model, Tikhonov regularization, restricted SVD, discrete Picard condition.

Part of the work of the first author was done while he was visiting Konrad-Zuse-Zentrum für Informationstechnik, Berlin.

The second author was supported by the Danish Natural Science Foundation, by the National Science Foundation under contract NSF-DMS87-14612, and by the Army Research Office under contract No. DAAL03-88-K-0085, during a visit to the Department of Mathematics, UCLA.

First, a word about our notation: $\|\cdot\|$ denotes the matrix and vector 2-norm, I_p is the identity matrix of order p , and A^+ denotes the pseudoinverse of A .

Let us consider the sensitivity of the solution to (1) to perturbations of the right-hand side \mathbf{b} . Let \mathbf{e} denote the perturbation, and let $\tilde{\mathbf{x}}$ denote the perturbed solution. Then the following approximate error bound follows from [10, equation (46)]:

$$(2) \quad \|\mathbf{x} - \tilde{\mathbf{x}}\| \lesssim \|A^+\| (1 + \|B\| \| (Q^T B)^+ \|) \|\mathbf{e}\|,$$

where the columns of Q form an orthonormal basis for the null space of A^T . We immediately see that if A is ill-conditioned, then \mathbf{x} may be very sensitive to perturbations. This is also clear from the analysis in [12], since we can always expect difficulties when dividing by the small generalized singular values of (A, B) . Equation (2) shows that a large $\| (Q^T B)^+ \|$ also indicates trouble.

In this paper we investigate the case where A is ill-conditioned while B is well-conditioned. To overcome the problems associated with the ill-conditioned A , we suggest adding Tikhonov regularization to problem (1) (Tikhonov regularization is discussed, e.g., in [3; 1, §26]). Thus, we propose the following *regularized Gauss-Markov problem*:

$$(3) \quad \min\{\|\mathbf{u}\|^2 + \lambda^2 \|C\mathbf{x}\|^2\} \quad \text{subject to} \quad A\mathbf{x} + B\mathbf{u} = \mathbf{b}.$$

Here, for simplicity, we assume that B and C have full rank,

$$(4a) \quad \begin{aligned} A &\in \mathbb{R}^{m \times n}, & B &\in \mathbb{R}^{m \times p}, & C &\in \mathbb{R}^{q \times n}, \\ \text{rank}(B) &= p \leq m, & \text{rank}(C) &= q \leq n \leq m. \end{aligned}$$

We also assume that

$$(4b) \quad \text{rank} \begin{pmatrix} A \\ C \end{pmatrix} = n,$$

which guarantees that the *regularized solution* \mathbf{x}_λ to (3) is unique for any $\lambda > 0$. Notice that we make no assumption about the rank of A , since this is not important in connection with Tikhonov regularization (cf., e.g., [4]). Typically, we will take C to be the identity matrix I_n or a well-conditioned discrete approximation to some derivative operator to ensure that the solution \mathbf{x}_λ is sufficiently “smooth”. The quantity λ is the *regularization parameter*, which controls the weight given to minimization of $\|C\mathbf{x}\|$ relative to minimization of $\|\mathbf{u}\|$.

We know that the regularized solution \mathbf{x}_λ to (3) is no longer an unbiased estimator (which is in fact the case for any regularized solution). However, inspired by the success of adding regularization to ill-conditioned least squares problems, we feel that \mathbf{x}_λ has other nice properties (cf. §3) that make it useful in connection with general Gauss-Markov linear models with ill-conditioned coefficient matrix A .

The paper is organized as follows. In §2 we introduce the restricted SVD and apply it as a tool for analyzing the model (3). In §3 we use these results to

describe the properties of the regularized solution \mathbf{x}_λ . In §4 we briefly discuss the discrete Picard condition as it applies to the regularized problem (3). Finally, in §5 we present a numerically stable algorithm for solving (3) efficiently.

2. AN RSVD ANALYSIS OF THE REGULARIZED GAUSS-MARKOV PROBLEM

We notice first that if $p < m$, then $\min\{\|B^+(A\mathbf{x} - \mathbf{b})\|^2 + \lambda^2\|C\mathbf{x}\|^2\}$ is not a valid formulation of (3), and we can therefore not base our analysis of (3) on Van Loan's S, T -singular values [13]. The proper tool to analyze (3) is the restricted SVD (RSVD) of (A, B, C) due to Zha [14]:

Theorem 1. *Let A, B , and C satisfy the assumptions in (4a) and (4b). Then there exist nonsingular matrices $X \in \mathbb{R}^{n \times n}$ and $Z \in \mathbb{R}^{m \times m}$, and orthogonal matrices $U \in \mathbb{R}^{p \times p}$ and $V \in \mathbb{R}^{q \times q}$ such that*

$$(5) \quad Z^T A X = \Sigma, \quad Z^T B U = M, \quad V^T C X = N,$$

where Σ, M , and N are pseudodiagonal matrices with nonnegative elements having the following structure:

$$(6) \quad \Sigma = \begin{bmatrix} \Sigma_A & 0 & 0 & 0 \\ 0 & I_j & 0 & 0 \\ 0 & 0 & I_k & 0 \\ 0 & 0 & 0 & I_l \\ 0 & 0 & 0 & 0 \end{bmatrix} \begin{matrix} s \\ j \\ k \\ l \\ u \end{matrix},$$

$$M = \begin{bmatrix} I_s & 0 \\ 0 & 0 \\ 0 & I_k \\ 0 & 0 \\ 0 & 0 \end{bmatrix} \begin{matrix} s \\ j \\ k \\ l \\ u \end{matrix}, \quad N = \begin{bmatrix} I_t & 0 & 0 & 0 \\ 0 & I_j & 0 & 0 \\ & t & j & k \\ & & & l \end{bmatrix} \begin{matrix} t \\ j \\ \\ \end{matrix}$$

and where

$$(6a) \quad \Sigma_A = \text{diag}(\sigma_1, \dots, \sigma_r) \in \mathbb{R}^{s \times t}, \quad \sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_r \geq 0, \\ r = \min\{s, t\}.$$

The dimensions of the submatrices are

$$(7) \quad j = r_{ab} + q - r_{abc}, \quad k = n + p - r_{abc}, \quad l = r_{abc} - p - q, \\ s = r_{abc} - n, \quad t = r_{abc} - r_{ab}, \quad u = m - r_{ab},$$

where $r_{ab} = \text{rank}(A, B)$ and $r_{abc} = \text{rank}\begin{pmatrix} A & B \\ C & 0 \end{pmatrix}$.

Proof. The proof of the RSVD as well as the notation is from [14, Theorem 4.2] with the simplifications imposed by our assumptions in (4a) and (4b). \square

Remark. In [5] it is shown that if A is ill-conditioned and C is well-conditioned, and if $A = \widehat{U}\widehat{\Sigma}\widehat{X}^{-1}$, $C = \widehat{V}\widehat{M}\widehat{X}^{-1}$ is the generalized SVD of (A, C) , then $\widehat{\Sigma}$ is ill-conditioned while \widehat{M} and \widehat{X} are well-conditioned. Using this result in the constructive proof for the RSVD [14], which consists of a sequence

of generalized SVD's, one can show that if A is ill-conditioned and both B and C are well-conditioned, then Σ is ill-conditioned while X and Z are well-conditioned.

Inserting the RSVD into (3) and using the fact that the 2-norm is invariant under orthogonal transformations, we immediately obtain the equivalent problem

$$(8) \quad \min\{\|\bar{\mathbf{u}}\|^2 + \lambda^2\|N\bar{\mathbf{x}}\|^2\} \quad \text{subject to} \quad \Sigma\bar{\mathbf{x}} + M\bar{\mathbf{u}} = \bar{\mathbf{b}},$$

where we have defined the transformed vectors $\bar{\mathbf{u}} = U^T \mathbf{u}$, $\bar{\mathbf{x}} = X^{-1} \mathbf{x}$, and $\bar{\mathbf{b}} = Z^T \mathbf{b}$. At this point, it is convenient to partition the matrices X and Z columnwise and to partition the vectors $\bar{\mathbf{x}}$, $\bar{\mathbf{u}}$, and $\bar{\mathbf{b}}$ elementwise according to the partitioning in (6),

$$(9a) \quad X = [X_t, X_j, X_k, X_l], \quad Z = [Z_s, Z_j, Z_k, Z_l, Z_u],$$

$$(9b) \quad \bar{\mathbf{x}} = \begin{bmatrix} \bar{\mathbf{x}}_t \\ \bar{\mathbf{x}}_j \\ \bar{\mathbf{x}}_k \\ \bar{\mathbf{x}}_l \end{bmatrix}, \quad \bar{\mathbf{u}} = \begin{bmatrix} \bar{\mathbf{u}}_s \\ \bar{\mathbf{u}}_k \end{bmatrix}, \quad \bar{\mathbf{b}} = \begin{bmatrix} \bar{\mathbf{b}}_s \\ \bar{\mathbf{b}}_j \\ \bar{\mathbf{b}}_k \\ \bar{\mathbf{b}}_l \\ \bar{\mathbf{b}}_u \end{bmatrix}.$$

The equality constraints in (8) then take the following simpler form:

$$(10) \quad \begin{aligned} \Sigma_A \bar{\mathbf{x}}_t + \bar{\mathbf{u}}_s &= \bar{\mathbf{b}}_s, & \bar{\mathbf{x}}_k + \bar{\mathbf{u}}_k &= \bar{\mathbf{b}}_k, & \mathbf{0} &= \bar{\mathbf{b}}_u. \\ \bar{\mathbf{x}}_j &= \bar{\mathbf{b}}_j, & \bar{\mathbf{x}}_l &= \bar{\mathbf{b}}_l, & & \end{aligned}$$

We immediately see that *consistency* of the model (3), i.e., the requirement that the right-hand side \mathbf{b} belongs to the range of the matrix (A, B) , corresponds to requiring that $\bar{\mathbf{b}}_u = \mathbf{0}$. Using the results in (10), the minimization problem in (8) can now be written as

$$(11) \quad \min \left\{ \left\| \begin{bmatrix} \bar{\mathbf{b}}_s \\ \bar{\mathbf{b}}_k \end{bmatrix} - \begin{bmatrix} \Sigma_A & 0 \\ 0 & I_k \end{bmatrix} \begin{bmatrix} \bar{\mathbf{x}}_t \\ \bar{\mathbf{x}}_k \end{bmatrix} \right\|^2 + \lambda^2 \left\| \begin{bmatrix} \bar{\mathbf{x}}_t \\ \bar{\mathbf{b}}_j \end{bmatrix} \right\|^2 \right\}.$$

The minimum is obtained for $\bar{\mathbf{x}}_k = \bar{\mathbf{b}}_k$, and $\bar{\mathbf{x}}_t$ is the solution to

$$(12) \quad \min\{\|\Sigma_A \bar{\mathbf{x}}_t - \bar{\mathbf{b}}_s\|^2 + \lambda^2 \|\bar{\mathbf{x}}_t\|^2\},$$

which is a discrete regularization problem in standard form and with the unique solution given by

$$(13a) \quad \bar{\mathbf{x}}_t = F_\lambda \Sigma_A^+ \bar{\mathbf{b}}_s,$$

where we have defined $F_\lambda = \text{diag}(f_i) \in \mathbb{R}^{t \times t}$ with diagonal elements

$$(13b) \quad f_i = \frac{\sigma_i^2}{\sigma_i^2 + \lambda^2}, \quad i = 1, \dots, t.$$

For more details about discrete standard-form regularization, cf., e.g., [4]. Notice that all parts of the solution $\bar{\mathbf{x}}$ are determined from (10) and (11). If

the matrix $\begin{pmatrix} A \\ C \end{pmatrix}$ does not have full rank, then there will also be a nonestimable (arbitrary) part of the solution.

The solution \mathbf{x}_λ to (3), as a function of λ , can thus be written as

$$(14) \quad \mathbf{x}_\lambda = X \begin{bmatrix} F_\lambda \Sigma_A^+ & 0 & 0 & 0 & 0 \\ 0 & I_j & 0 & 0 & 0 \\ 0 & 0 & I_k & 0 & 0 \\ 0 & 0 & 0 & I_l & 0 \end{bmatrix} Z^T \mathbf{b}$$

$$= X_t F_\lambda \Sigma_A^+ Z_s^T \mathbf{b} + X_j Z_j^T \mathbf{b} + X_k Z_k^T \mathbf{b} + X_l Z_l^T \mathbf{b}.$$

In particular, if $q = n$ then $k + l = 0$, so that the last two terms in (14) vanish. In general, we have $\mathbf{x}_\lambda \rightarrow \mathbf{x}_0 = X \Sigma^+ Z^T \mathbf{b}$ for $\lambda \rightarrow 0$. Note that the matrix $X \Sigma^+ Z^T$ is not a weighted pseudoinverse of A as defined by Eldén [2].

3. SOME PROPERTIES OF THE REGULARIZED SOLUTION

We shall now describe some of the nice features of the regularized solution \mathbf{x}_λ to (3). First of all, we see that if A has any small σ_i , reflecting the ill-conditioning of A , then the norm of the unregularized solution $\mathbf{x}_0 = X \Sigma^+ Z^T \mathbf{b}$ may be very large because of the division by these small σ_i . For the same reason, \mathbf{x}_0 is very sensitive to perturbations of \mathbf{b} . Consider now the first term in the expression (14) for the regularized solution \mathbf{x}_λ :

$$X_t F_\lambda \Sigma_A^+ Z_s^T \mathbf{b} = \sum_{i=1}^r \frac{\sigma_i^2}{\sigma_i^2 + \lambda^2} \frac{z_i^T \mathbf{b}}{\sigma_i} \mathbf{x}_i.$$

This is the only term where small σ_i occur. We immediately see from this expression that by choosing a suitable regularization parameter λ somewhere between σ_r and σ_1 , we are able to ‘filter out’ the contributions to \mathbf{x}_λ corresponding to the small σ_i via the matrix F_λ . In this way, we can use λ to control the norm and the sensitivity of \mathbf{x}_λ at the expense of neglecting a (small) part of the information in the right-hand side \mathbf{b} . This is, in fact, completely analogous to regularization of least squares problems [4].

Next, we prove that the regularized problem (3) is indeed better conditioned than the original problem (1). For simplicity, we restrict the perturbations to the right-hand side.

Theorem 2. *Let the perturbed right-hand side be $\tilde{\mathbf{b}} = \mathbf{b} + \mathbf{e}$, let $\tilde{\mathbf{x}}_\lambda$ denote the corresponding perturbed solution, and let \mathbf{x}_0 and $\tilde{\mathbf{x}}_0$ denote the solutions for $\lambda = 0$. If $\sigma_r \leq \lambda \leq \sigma_1$, then*

$$(15) \quad \frac{\|\mathbf{x}_\lambda - \tilde{\mathbf{x}}_\lambda\|}{\|\mathbf{x}_\lambda\|} \leq \frac{\|\Sigma\|}{\min\{2\lambda, 1\}} \kappa(X) \kappa(Z) \frac{\|\mathbf{e}\|}{\|\mathbf{b}_\lambda\|},$$

$$(16) \quad \frac{\|\mathbf{x}_0 - \tilde{\mathbf{x}}_0\|}{\|\mathbf{x}_0\|} \leq \frac{\|\Sigma\|}{\sigma_r} \kappa(X) \kappa(Z) \frac{\|\mathbf{e}\|}{\|\mathbf{b}_0\|}.$$

Here, $\|\Sigma\| = \max\{\sigma_1, 1\}$, $\mathbf{b}_\lambda = A \mathbf{x}_\lambda$, $\mathbf{b}_0 = A \mathbf{x}_0$, and $\kappa(X)$ and $\kappa(Z)$ are the condition numbers of X and Z , respectively.

Proof. The relations $\|\mathbf{b}_\lambda\| = \|\mathbf{A}\mathbf{x}_\lambda\| \leq \|Z^{-1}\| \|\Sigma\| \|X^{-1}\| \|\mathbf{x}_\lambda\|$ and $\|\mathbf{x}_\lambda - \tilde{\mathbf{x}}_\lambda\| \leq \|X\| \max\{\|F_\lambda \Sigma_A^+\|, 1\} \|Z\| \|\mathbf{e}\|$ give

$$\frac{\|\mathbf{x}_\lambda - \tilde{\mathbf{x}}_\lambda\|}{\|\mathbf{x}_\lambda\|} \leq \kappa(X) \|\Sigma\| \max\{\|F_\lambda \Sigma_A^+\|, 1\} \kappa(Z) \frac{\|\mathbf{e}\|}{\|\mathbf{b}_\lambda\|}.$$

It is straightforward to show that

$$\|F_\lambda \Sigma_A^+\| = \max\{f_i/\sigma_i\} = \max\{\sigma_i/(\sigma^2 + \lambda^2)\} \leq 1/(2\lambda),$$

so that $\max\{\|F_\lambda \Sigma_A^+\|, 1\} = 1/\min\{2\lambda, 1\}$. Since $\|\Sigma\| = \max\{\|\Sigma_A\|, 1\} = \max\{\sigma_1, 1\}$, this yields (15). Equation (16) is derived analogously by using that $\|\mathbf{x}_0 - \tilde{\mathbf{x}}_0\| \leq \|X\| \|\Sigma^+\| \|Z\| \|\mathbf{e}\|$ and that $\|\Sigma^+\| = \sigma_r^{-1}$. \square

Remark. Theorem 2 shows that the condition number κ_λ associated with (3) satisfies

$$(17) \quad \kappa_\lambda \equiv \limsup_{\|\mathbf{e}\| \rightarrow 0} \frac{\|\mathbf{x}_\lambda - \tilde{\mathbf{x}}_\lambda\|}{\|\mathbf{x}_\lambda\|} \leq \frac{\|\Sigma\|}{\min\{2\lambda, 1\}} \kappa(X) \kappa(Z).$$

The key point here is that Theorem 2 shows that it is always possible to choose λ such that \mathbf{x}_λ is much less sensitive to perturbations than \mathbf{x}_0 . Thus, we can say that for appropriate regularization parameter λ , (3) is better conditioned than (1).

Another important property of introducing regularization in (3) is that it decreases the variance of the solution \mathbf{x}_λ , compared to the variance of the solution \mathbf{x}_0 to (1) without regularization. Since the variance-covariance matrix associated with $\bar{\mathbf{u}} = U^T \mathbf{u}$ is $V(\bar{\mathbf{u}}) = s^2 I_p$, it is easy to show that the variance-covariance matrix $V(\bar{\mathbf{x}}_r)$ associated with the regularized solution vector $\bar{\mathbf{x}}_r$ is

$$(18) \quad V(\bar{\mathbf{x}}_r) = s^2 (F_\lambda \Sigma_A^+)^2.$$

We readily see that if λ is chosen suitably, somewhere between σ_r and σ_1 , then the elements of this matrix are numerically much smaller than those of the variance-covariance matrix $s^2 (\Sigma_A^+)^2$ corresponding to $\lambda = 0$.

In this discussion we have not considered the “smoothness” of \mathbf{x}_λ . We felt that such an analysis can be performed in analogy with that in [5]. For example, we know that the null space of C , which is spanned by the columns of X_k and X_l , is always “smooth” (in the sense of few zero crossings) when C is a discrete approximation to a derivative operator—thus ensuring that the component $X_k Z_k^T \mathbf{b} + X_l Z_l^T \mathbf{b}$ in \mathbf{x}_λ (14) is also “smooth”. However, we were not able to derive any results about the “smoothness” of the columns of the submatrices X_l and X_j .

4. THE DISCRETE PICARD CONDITION

Of course, the introduction of regularization in (3) changes the solution \mathbf{x}_λ compared to the unregularized solution to (1). The purpose of this section is to investigate the difference between these solutions. In this connection, notice that

if A does not have full rank, then the solution to (1) is not unique: the general solution can always be written as the estimable part of the solution plus an arbitrary amount of the nonestimable part of the solution [12]. For $\lambda \rightarrow 0$, the regularized solution \mathbf{x}_λ converges to a member of this general solution (but not necessarily to the estimable part; we can only guarantee this if $\text{rank}(A) = n$). It is therefore correct to compare the regularized solution \mathbf{x}_λ to the solution $\mathbf{x}_0 = X\Sigma^+Z^T\mathbf{b}$ obtained from (14) by setting $\lambda = 0$.

An analysis of the regularization error $\mathbf{x}_0 - \mathbf{x}_\lambda$ for general A, B, C , and \mathbf{b} is probably not possible. Instead, we use the same technique as in [4, 6]: we assume a very simple (but still realistic) “model” of the right-hand side and determine the conditions under which the regularization error is guaranteed to be small. Our “model” here is

$$(19) \quad z_i^T \mathbf{b} = \begin{cases} \sigma_i^\alpha, & i = 1, \dots, r, \\ \sigma_r^\alpha, & i = r + 1, \dots, s, \end{cases} \quad \alpha \geq 0.$$

The parameter α controls the decay of the $z_i^T \mathbf{b}$ relative to the decay of the corresponding σ_i , in such a way that the $z_i^T \mathbf{b}$ decay faster to zero than the σ_i for $\alpha > 1$. A direct analysis of $\mathbf{x}_0 - \mathbf{x}_\lambda$ is very difficult, so we multiply by the well-conditioned matrix X^{-1} and consider instead $X^{-1}(\mathbf{x}_0 - \mathbf{x}_\lambda)$:

Theorem 3. Let $\mathbf{x}_0 = X\Sigma^+Z^T\mathbf{b}$ be defined as the solution \mathbf{x}_λ (14) with $\lambda = 0$, and let the right-hand side \mathbf{b} satisfy (19). If $\sigma_r \leq \lambda \leq \sigma_1$, then

$$(20) \quad \frac{\|X^{-1}(\mathbf{x}_0 - \mathbf{x}_\lambda)\|}{\|X^{-1}\mathbf{x}_0\|} \leq \begin{cases} \sqrt{r}(\sigma_1/\sigma_r)^{1-\alpha}, & 0 \leq \alpha < 1, \\ \sqrt{r}(\lambda/\sigma_1)^{\alpha-1}, & 1 \leq \alpha < 3, \\ \sqrt{r}(\lambda/\sigma_1)^2, & \alpha \geq 3. \end{cases}$$

Proof. We have

$$\begin{aligned} \|X^{-1}(\mathbf{x}_0 - \mathbf{x}_\lambda)\| &= \|(\Sigma_A^+ - F_\lambda \Sigma_A^+)Z_s^T \mathbf{b}\| = \|(I_t - F_\lambda)\Sigma_A^+Z_s^T \mathbf{b}\| \\ &\leq \sqrt{r}\|(I_t - F_\lambda)\Sigma_A^+Z_s^T \mathbf{b}\|_\infty = \sqrt{r} \max\{(1 - f_i)\sigma_i^{\alpha-1}\}. \end{aligned}$$

Here, $(1 - f_i)\sigma_i^{\alpha-1} = \lambda^2(\sigma_i^2 + \lambda^2)^{-1}\sigma_i^{\alpha-1} = \lambda^2\phi(\sigma_i)$, where we have defined $\phi(\sigma) \equiv \sigma^{\alpha-1}/(\sigma^2 + \lambda^2)$. It is easy to show the following: For $0 \leq \alpha < 1$: $\phi(\sigma)$ is decreasing, so that

$$\lambda^2\phi(\sigma) \leq \lambda^2\phi(\sigma_r) = \frac{\lambda^2}{\sigma_r^2 + \lambda^2}\sigma_r^{\alpha-1} \leq \sigma_r^{\alpha-1};$$

for $1 \leq \alpha < 3$: $\phi(\sigma)$ has its maximum at $\sigma^2 = \lambda^2(\alpha - 1)/(3 - \alpha)$, and

$$\begin{aligned} \lambda^2\phi(\sigma) &\leq \frac{\lambda^2}{\lambda^2(\alpha - 1)/(3 - \alpha) + \lambda^2} \left(\frac{\alpha - 1}{3 - \alpha}\lambda^2\right)^{(\alpha-1)/2} \\ &= \frac{3 - \alpha}{2} \left(\frac{\alpha - 1}{3 - \alpha}\right)^{(\alpha-1)/2} \lambda^{\alpha-1} \\ &= \frac{1}{2}(\alpha - 1)^{(\alpha-1)/2}(3 - \alpha)^{(3-\alpha)/2}\lambda^{\alpha-1} \leq \lambda^{\alpha-1}; \end{aligned}$$

for $\alpha \geq 3$: $\phi(\sigma)$ is increasing, so that

$$\lambda^2 \phi(\sigma) \leq \lambda^2 \phi(\sigma_1) = \frac{\lambda^2}{\sigma_1^2 + \lambda^2} \sigma_1^{\alpha-1} \leq \lambda^2 \sigma_1^{\alpha-3}.$$

Now, let Σ_0 denote Σ (6a) with the three identity matrices replaced by 0. Then

$$\|Z_s^T \mathbf{b}\| = \|\Sigma_0 X^{-1} \mathbf{x}_0\| \leq \|\Sigma_A\| \|X^{-1} \mathbf{x}_0\| = \sigma_1 \|X^{-1} \mathbf{x}_0\|.$$

From the definition (19) of \mathbf{b} we also have $\|Z_s^T \mathbf{b}\| \geq \|Z_s^T \mathbf{b}\|_\infty = \sigma_1^\alpha$. Thus, $1/\|X^{-1} \mathbf{x}_0\| \leq 1/\sigma_1^{\alpha-1}$. Together, these formulae lead to (20). \square

Not surprisingly, we see that in order to ensure a small upper bound for the regularization error, we must require that the coefficients $|z_i^T \mathbf{b}|$ decay to zero *faster* than the σ_i . We also see that the faster the decay, the better \mathbf{x}_λ approximates \mathbf{x}_0 . Following an idea in [6], we are then led to the following definition of the *discrete Picard condition* for the regularized Gauss-Markov problem (3):

Definition 4. The discrete Picard condition (DPC). The right-hand side \mathbf{b} in (3) satisfies the DPC if, for all numerically nonzero σ_i , the coefficients $|z_i^T \mathbf{b}|$ in average decay to zero faster than the σ_i .

If the underlying, unperturbed right-hand side in (3) does not satisfy the DPC, then there is no point in trying to solve (3) at all, because \mathbf{x}_λ does not approximate the true solution \mathbf{x}_0 for any value of λ . If, on the other hand, the unperturbed right-hand side \mathbf{b} satisfies the DPC, and if the given $\tilde{\mathbf{b}} = \mathbf{b} + \mathbf{e}$ (which is contaminated with errors) is not completely dominated by the errors \mathbf{e} , then $\tilde{\mathbf{b}}$ actually satisfies the DPC for $i \leq K$, where K is determined by the magnitude and the statistical distribution of the errors. Hence, if we choose $\lambda \approx \sigma_K$, then the effect of regularization is to dampen the contributions to \mathbf{x}_λ corresponding to the small $\sigma_i < \lambda$. In other words, we can regard the addition of regularization to the linear model as a means for producing a slightly perturbed model that is guaranteed to satisfy the DPC, thus ensuring that the regularized solution \mathbf{x}_λ is a meaningful estimator. For more details, and how to implement a check for satisfaction of the DPC in practice, cf. [6].

As in any regularization problem, the practical question of choosing a suitable regularization parameter λ is difficult to answer in general terms, essentially because the (restricted) singular values are not available [3, §§3 and 4]. In fact, no perfectly general algorithm or principle seems to exist, and often an interactive approach using plots of the solution \mathbf{x} , the derivative $C\mathbf{x}$, and the residual vector $\mathbf{r} = A\mathbf{x} - \mathbf{b}$, as well as their norms, is the best approach. Recent results in connection with discrete regularization do, however, suggest a near relationship between such an intuitive interactive approach and the method of *generalized cross-validation*. The basic idea is to plot the solution norm versus the residual norm as a parametrized curve with λ as the parameter, and this

curve will typically have a more or less distinct L -shaped corner at the optimal value of λ . (See [4, §5] and, in particular, [5, §5] for more details.)

In connection with the regularized Gauss-Markov problem (3), our analysis here has shown that the basic problem to solve is the discrete regularization problem (12). Therefore, we would ideally like to consider a plot of $\|\bar{x}_t\|$ versus $\|\Sigma_A \bar{x}_t - \bar{b}_s\|$; but these quantities are not readily available. However, it is easy to see that $\|C\mathbf{x}\|^2 = \|\bar{x}_t\|^2 + \|\bar{x}_j\|^2$ and (since B has full rank)

$$\|B^+ \mathbf{r}\|^2 = \|B^+(A\mathbf{x} - \mathbf{b})\|^2 = \|\mathbf{u}\|^2 = \|\bar{\mathbf{u}}\|^2 = \|\bar{\mathbf{u}}_s\|^2 + \|\bar{\mathbf{u}}_k\|^2;$$

i.e., the curve of $\|C\mathbf{x}\|$ versus $\|B^+ \mathbf{r}\|$ is merely a translation of the $(\|\Sigma_A \bar{x}_t - \bar{b}_s\|, \|\bar{x}_t\|)$ -curve. Hence, the choice of λ might as well be based on the plot of the $(\|B^+ \mathbf{r}\|, \|C\mathbf{x}\|)$ -curve, using exactly the same ideas as described in [4, 5]. In this connection, note that $\|B^+ \mathbf{r}\|$ is easily computed via a QR factorization of B (which has to be computed anyway; cf. Step 1 in the next section).

5. A NUMERICAL ALGORITHM

In this section we describe an algorithm for computing the unique regularized solution \mathbf{x}_λ to (3). It is easy to see that (3) can be reformulated as

$$\min \left\| \begin{bmatrix} \lambda C & 0 \\ 0 & I_p \end{bmatrix} \begin{bmatrix} \mathbf{x} \\ \mathbf{u} \end{bmatrix} \right\| \quad \text{subject to} \quad [A, B] \begin{bmatrix} \mathbf{x} \\ \mathbf{u} \end{bmatrix} = \mathbf{b},$$

which is a simple equality-constrained linear least squares problem. Algorithms such as those described in [8], especially the null-space method [8, Chapter 20], can be applied directly to solve the above problem. However, as is also pointed out in [10] (which is in the setting of general Gauss-Markov linear models without regularization), such an approach does not treat \mathbf{x} , \mathbf{u} , A , B , and C separately and in turn cannot take advantage of any special structure of the problem. Even worse is the case when solutions are required for several different λ , in which case one has to solve a new equality-constrained linear least squares problem from the beginning for each λ . The following algorithm tries to take these aspects into account, and is inspired by the work of Paige [10, 11]. The first three steps of our algorithm can be considered as a way to transform the problem to a simpler form, only using the data matrices A , B , and C . When a new λ is chosen, one only needs to restart the algorithm from Step 4. The algorithm requires that assumptions (4a) and (4b) be satisfied.

Step 1. Make a QR decomposition of B , so that $B = Q \begin{bmatrix} B_1 \\ 0 \end{bmatrix}$, where $B_1 \in \mathbb{R}^{p \times p}$ is upper triangular and nonsingular. Let

$$(21) \quad Q^T [A, b] = \begin{bmatrix} A_1 & \mathbf{b}_1 \\ A_2 & \mathbf{b}_2 \end{bmatrix} \begin{matrix} p \\ m - p \end{matrix}.$$

Step 2. Make the following decomposition of A_2 :

$$A_2 U = \begin{bmatrix} 0, & A_{22} \end{bmatrix} \begin{matrix} m - p, \\ n - i \quad i \end{matrix}$$

so that U is orthogonal and $A_{22} \in \mathbb{R}^{(m-p) \times i}$ is of full column rank i . Let

$$(22) \quad \begin{bmatrix} A_1 \\ C \end{bmatrix} U = \begin{bmatrix} A_{11} & A_{12} \\ C_1 & C_2 \\ n-i & i \end{bmatrix} \quad \text{and} \quad U^T x = \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} \begin{matrix} n-i \\ i \end{matrix}.$$

Then $\begin{bmatrix} A_{11} \\ C_1 \end{bmatrix}$ is of full column rank.

Step 3 (including consistency check). Make a QR decomposition of A_{22} , so that $A_{22} = Q_1 \begin{bmatrix} \tilde{A}_{22} \\ 0 \end{bmatrix}$, where $\tilde{A}_{22} \in \mathbb{R}^{i \times i}$ is upper triangular and nonsingular, and let

$$Q_1^T \mathbf{b}_2 = \begin{bmatrix} \mathbf{b}_2^{(1)} \\ \mathbf{b}_2^{(2)} \end{bmatrix}, \quad \text{with } \mathbf{b}_2^{(1)} \in \mathbb{R}^i.$$

Then the regularized general Gauss-Markov linear model is consistent only if $\mathbf{b}_2^{(2)} = \mathbf{0}$. In this case,

$$(23) \quad \mathbf{x}_2 = \tilde{A}_{22}^{-1} \mathbf{b}_2^{(1)}.$$

After some manipulation, we obtain the following ordinary least squares problem, which only involves the component \mathbf{x}_1 :

$$(24) \quad \min \left\| \begin{bmatrix} \lambda C_1 \\ B_1^{-1} A_{11} \end{bmatrix} \mathbf{x}_1 - \begin{bmatrix} \lambda C_2 \mathbf{x}_2 \\ B_1^{-1} (\mathbf{b}_1 - A_{12} \mathbf{x}_2) \end{bmatrix} \right\|.$$

Step 4. The observation here is that we do not need to explicitly compute the matrix products involving B_1^{-1} in order to solve the above problem, and this enhances the numerical stability of the algorithm. In fact, (24) is equivalent to the “standard” problem described in [11] with

$$\tilde{\mathbf{y}} = \begin{bmatrix} \lambda C_2 \mathbf{x}_2 \\ \mathbf{b}_1 - A_{12} \mathbf{x}_2 \end{bmatrix}, \quad \tilde{C} = \begin{bmatrix} \lambda C_1 \\ A_{11} \end{bmatrix}, \quad \text{and} \quad \tilde{B} = \begin{bmatrix} I & 0 \\ 0 & B_1 \end{bmatrix}.$$

We only need to choose orthogonal \tilde{P} and \tilde{Q} such that

$$(25) \quad \begin{aligned} & \tilde{Q}^T [\tilde{C}, \tilde{\mathbf{y}}, \tilde{B}] \begin{bmatrix} I & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & \tilde{P} \end{bmatrix} \\ & = \begin{bmatrix} R & \tau & R_1 & \mathbf{r} & R_{12} \\ 0 & \eta & 0 & \rho & \mathbf{w}^T \\ 0 & 0 & 0 & 0 & R_2 \end{bmatrix} \begin{matrix} n-i \\ 1 \\ n-i & 1 \end{matrix}, \end{aligned}$$

where R , R_1 , and R_2 are upper triangular and nonsingular. Now, if we let

$$\tilde{P}^T \mathbf{u} = \mathbf{v} = \begin{bmatrix} \mathbf{v}_1 \\ \mu \\ \mathbf{v}_2 \end{bmatrix} \begin{matrix} n-i \\ 1 \\ p+q-(n-i)-1 \end{matrix},$$

then \mathbf{x}_1 can be obtained by solving the following nonsingular upper triangular linear system:

$$(26) \quad \begin{bmatrix} R & \mathbf{r} \\ \mathbf{0}^T & \rho \end{bmatrix} \begin{bmatrix} \mathbf{x}_1 \\ \mu \end{bmatrix} = \begin{bmatrix} \tau \\ \eta \end{bmatrix}.$$

For details, see [11].

Step 5. The unique solution is given by

$$(27) \quad \mathbf{x}_\lambda = U \begin{bmatrix} \mathbf{x}_1 \\ \mathbf{x}_2 \end{bmatrix}.$$

We now give more details about the factorization of (25). Since \tilde{B} is block diagonal, we can exploit this in the reduction step. In order to illustrate the situation, we give below a low-dimensional example with $p = q = n - i + 1 = 3$. We first use a QR decomposition to transform $[\lambda C_1, \lambda C_2 \mathbf{x}_2]$ to upper triangular form. Since the first block of \tilde{B} is an identity matrix, the effect of this transformation can be compensated by multiplying \mathbf{u} by the orthogonal matrix, so that at the initial stage we have

$$([\tilde{C}, \tilde{\mathbf{y}}], \tilde{B}) = \left(\begin{bmatrix} \times & \times & \times \\ & \times & \times \\ \times & \times & \times \\ \times & \times & \times \\ \times & \times & \times \end{bmatrix}, \begin{bmatrix} \times & & & & & & \\ & \times & & & & & \\ & & \times & & & & \\ & & & \times & & & \\ & & & & \times & \times & \times \\ & & & & & \times & \times \\ & & & & & & \times \end{bmatrix} \right).$$

We can now reduce the leftmost matrix $[\tilde{C}, \tilde{\mathbf{y}}]$ to triangular form, while maintaining the triangular form of the rightmost matrix \tilde{B} , by a sequence of left and right Givens transformations. Figure 1 below illustrates how to eliminate the first column of the lower part of the first matrix. The same procedure can be continued in a similar way until we obtain the decomposition in (25). Throughout, \rightarrow shows the two rows or columns involved in the Givens transformation, and \circ indicates the element being annihilated.




FIGURE 1

The first steps in reducing the leftmost matrix to triangular form while maintaining the triangular form of the rightmost matrix

ACKNOWLEDGMENTS

H. Zha wishes to thank Professor Dr. Deuffhard at the Konrad-Zuse-Zentrum for his support and encouragement. P. C. Hansen would like to thank Professor Tony Chan for the invitation to UCLA and for providing nice working conditions there.

BIBLIOGRAPHY

1. Å. Björck, *Least squares methods*, Handbook of Numerical Analysis, Vol. I: Finite Difference Methods—Solution of Equations in \mathbb{R}^n (P. G. Ciarlet and J. L. Lions, eds.), Elsevier, 1990.
2. L. Eldén, *A weighted pseudoinverse, generalized singular values, and constrained least squares problems*, BIT **2** (1982), 487–502.
3. C. W. Groetsch, *The theory of Tikhonov regularization for Fredholm equations of the first kind*, Research Notes in Math., vol. 105, Pitman, New York, 1984.
4. P. C. Hansen, *Truncated SVD solutions to discrete ill-posed problems with ill-determined numerical rank*, SIAM J. Sci. Statist. Comput. **11** (1990), 503–518.
5. —, *Regularization, GSVD and truncated GSVD*, BIT **29** (1989), 491–504.
6. —, *The discrete Picard condition for discrete ill-posed problems*, BIT **30** (1990), (to appear).
7. S. Kourouklis and C. C. Paige, *A constrained least squares approach to the general Gauss-Markov linear model*, J. Amer. Statist. Assoc. **76** (1981), 620–625.
8. C. L. Lawson and R. J. Hanson, *Solving least squares problems*, Prentice-Hall, Englewood Cliffs, N. J., 1974.
9. C. C. Paige, *Numerically stable computations for general univariate linear models*, Comm. Statist. B—Simulation Comput. **7** (1978), 437–453.
10. —, *Computer solution and perturbation analysis of generalized linear least squares problems*, Math. Comp. **33** (1979), 171–183.
11. —, *Fast numerically stable computations for generalized least squares problems*, SIAM J. Numer. Anal. **16** (1979), 165–171.
12. —, *The general linear model and the generalized singular value decomposition*, Linear Algebra Appl. **70** (1985), 269–284.
13. C. F. Van Loan, *Generalizing the singular value decomposition*, SIAM J. Numer. Anal. **13** (1976), 76–83.
14. H. Zha, *Restricted singular value decomposition of matrix triplets*, Report SC-89-2, Konrad-Zuse-Zentrum für Informationstechnik Berlin, 1989 (submitted to SIAM J. Matrix. Anal. Appl.).

SC/CM COMPUTER SCIENCE DEPARTMENT, STANFORD UNIVERSITY, STANFORD, CALIFORNIA 90305

E-mail address: na.zha@na-net.stanford.edu

UNI · C LYNGBY, BUILDING 305, TECHNICAL UNIVERSITY OF DENMARK, DK-2800 LYNGBY, DENMARK

E-mail address: na.phansen@na-net.stanford.edu