

PARAMETER ESTIMATION IN THE PRESENCE OF BOUNDED DATA UNCERTAINTIES*

S. CHANDRASEKARAN[†], G. H. GOLUB[‡], M. GU[§], AND A. H. SAYED[¶]

Abstract. We formulate and solve a new parameter estimation problem in the presence of data uncertainties. The new method is suitable when a priori bounds on the uncertain data are available, and its solution leads to more meaningful results, especially when compared with other methods such as total least-squares and robust estimation. Its superior performance is due to the fact that the new method guarantees that the effect of the uncertainties will never be unnecessarily overestimated, beyond what is reasonably assumed by the a priori bounds. A geometric interpretation of the solution is provided, along with a closed form expression for it. We also consider the case in which only selected columns of the coefficient matrix are subject to perturbations.

Key words. least-squares estimation, regularized least-squares, ridge regression, total least-squares, robust estimation, modeling errors, secular equation

AMS subject classifications. 15A06, 65F05, 65F10, 65F35, 65K10, 93C41, 93E10, 93E24

PII. S0895479896301674

1. Introduction. The central problem in estimation is to recover, to good accuracy, a set of unobservable parameters from corrupted data. Several optimization criteria have been used for estimation purposes over the years, but the most important, at least in the sense of having had the most applications, are criteria that are based on quadratic cost functions. The most striking among these is the linear least-squares criterion, which was first developed by Gauss (ca. 1795) in his work on celestial mechanics. Since then, it has enjoyed widespread popularity in many diverse areas as a result of its attractive computational and statistical properties (see, e.g., [4, 8, 10, 13]). Among these attractive properties, the most notable are the facts that least-squares solutions can be explicitly evaluated in closed forms, they can be recursively updated as more input data is made available, and they are also maximum likelihood estimators in the presence of normally distributed measurement noise.

Alternative optimization criteria, however, have been proposed over the years including, among others, regularized least-squares [4], ridge regression [4, 10], total least-squares [2, 3, 4, 7], and robust estimation [6, 9, 12, 14]. These different formulations allow, in one way or another, incorporation of further a priori information about the unknown parameter into the problem statement. They are also more effective in the presence of data errors and incomplete statistical information about the exogenous signals (or measurement errors).

Among the most notable variations is the total least-squares (TLS) method, also known as orthogonal regression or errors-in-variables method in statistics and system

*Received by the editors April 8, 1996; accepted for publication (in revised form) by N. J. Higham February 13, 1997.

<http://www.siam.org/journals/simax/19-1/30167.html>

[†]Department of Electrical and Computer Engineering, University of California, Santa Barbara, CA 93016 (shiv@ece.ucsb.edu).

[‡]Department of Computer Science, Stanford University, Stanford, CA 94305 (golub@scm.stanford.edu).

[§]Department of Mathematics, University of California, Los Angeles, CA 90095 (mgu@math.ucla.edu).

[¶]Department of Electrical Engineering, University of California, Los Angeles, CA 90095 (sayed@ee.ucla.edu). The work of this author was supported in part by National Science Foundation award MIP-9796147.

identification [11]. In contrast to the standard least-squares problem, the TLS formulation allows for errors in the data matrix. But it still exhibits certain drawbacks that degrade its performance in practical situations. In particular, it may unnecessarily overemphasize the effect of noise and uncertainties and can, therefore, lead to overly conservative results.

More specifically, assume $A \in \mathbf{R}^{m \times n}$ is a given full rank matrix with $m \geq n$, $b \in \mathbf{R}^m$ is a given vector, and consider the problem of solving the inconsistent linear system $A\hat{x} \approx b$ in the least-squares sense. The TLS solution assumes data uncertainties in A and proceeds to correct A and b by replacing them by their projections, \hat{A} and \hat{b} , onto a specific subspace and by solving the consistent linear system of equations $\hat{A}\hat{x} = \hat{b}$. The spectral norm of the correction $(A - \hat{A})$ in the TLS solution is bounded by the smallest singular value of $[A \ b]$. While this norm might be small for vectors b that are close enough to the range space of A , it need not always be so. In other words, the TLS solution may lead to situations in which the correction term is unnecessarily large.

Consider, for example, a situation in which the uncertainties in A are very small, say, A is almost known exactly. Assume further that b is far from the column space of A . In this case, it is not difficult to visualize that the TLS solution will need to rotate (A, b) into (\hat{A}, \hat{b}) and may therefore end up with an overly corrected approximant for A , despite the fact that A is almost exact.

These facts motivate us to introduce a new parameter estimation formulation with prior bounds on the size of the allowable corrections to the data. More specifically, we formulate and solve a new estimation problem that is more suitable for scenarios in which a priori bounds on the uncertain data are known. The solution leads to more meaningful results in the sense that it guarantees that the effect of the uncertainties will never be unnecessarily overestimated, beyond what is reasonably assumed by the a priori bounds.

We note that, while preparing this paper, the related work [1] has come to our attention, where the authors have independently formulated and solved a similar estimation problem by using (convex) semidefinite programming techniques and interior-point methods. The resulting computational complexity of the proposed solution is $O(nm^2 + m^{3.5})$, where n is the smaller matrix dimension.

The solution proposed in this paper proceeds by first providing a geometric formulation of the problem, followed by an algebraic derivation that establishes that the optimal solution can in fact be obtained by solving a related regularized problem. The parameter of the regularization step is further shown to be obtained as the unique positive root of a secular equation and as a function of the given data. In this sense, the new formulation turns out to provide automatic regularization and, hence, has some useful regularization properties: the regularization parameter is not selected by the user but rather determined by the algorithm. Our solution involves an SVD step, and its computational complexity amounts to $O(mn^2 + n^3)$, where n is again the smaller matrix dimension. A summary of the problem and its solution is provided in section 3.4. (Other problem formulations are studied in [15].)

2. Problem formulation. Let $A \in \mathbf{R}^{m \times n}$ be a given matrix with $m \geq n$ and $b \in \mathbf{R}^m$ a given vector, both of which are assumed to be linearly related via an unknown vector of parameters $x \in \mathbf{R}^n$,

$$(2.1) \quad b = Ax + v .$$

The vector $v \in \mathbf{R}^m$ denotes measurement noise and it explains the mismatch between Ax and the given vector (or observation) b .

We assume that the “true” coefficient matrix is $A + \delta A$ and that we only know an upper bound on the 2-induced norm of the perturbation δA ,

$$(2.2) \quad \|\delta A\|_2 \leq \eta ,$$

with η being known. Likewise, we assume that the “true” observation vector is $b + \delta b$ and that we know an upper bound η_b on the Euclidean norm of the perturbation δb ,

$$(2.3) \quad \|\delta b\|_2 \leq \eta_b .$$

We then pose the problem of finding an estimate that performs “well” for any allowed perturbation $(\delta A, \delta b)$. More specifically, we pose the following min-max problem.

PROBLEM 1. *Given $A \in \mathbf{R}^{m \times n}$, with $m \geq n$, $b \in \mathbf{R}^m$, and nonnegative real numbers (η, η_b) , determine, if possible, an \hat{x} that solves*

$$(2.4) \quad \min_{\hat{x}} \max \{ \| (A + \delta A) \hat{x} - (b + \delta b) \|_2 : \|\delta A\|_2 \leq \eta, \|\delta b\|_2 \leq \eta_b \} .$$

The situation is depicted in Fig. 2.1. Any particular choice for \hat{x} would lead to many residual norms,

$$\| (A + \delta A) \hat{x} - (b + \delta b) \|_2 ,$$

one for each possible choice of A in the disc $(A + \delta A)$ and b in the disc $(b + \delta b)$. A second choice for \hat{x} would lead to other residual norms, the maximum value of which need not be the same as the first choice. We want to choose an estimate \hat{x} that minimizes the maximum possible residual norm. This is depicted in Fig. 2.2 for two choices, say \hat{x}_1 and \hat{x}_2 . The curves show the values of the residual norms as a function of $(A + \delta A, b + \delta b)$.

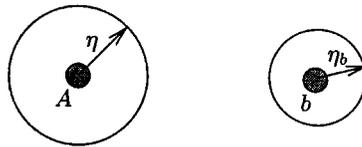


FIG. 2.1. *Geometric interpretation of the new least-squares formulation.*

We note that if $\eta = 0 = \eta_b$, then problem (2.4) reduces to a standard least-squares problem. Therefore we shall assume throughout that $\eta > 0$. (It will turn out that the solution to the above min-max problem is independent of η_b .)

2.1. A geometric interpretation. The min-max problem admits an interesting geometric formulation that highlights some of the issues involved in its solution.

For this purpose, and for the sake of illustration, assume we have a unit-norm vector b , $\|b\|_2 = 1$, with no uncertainties in it ($\eta_b = 0$). Assume further that A is simply a column vector, say a , with $\eta \neq 0$. That is, only A is assumed to be uncertain with perturbations that are bounded by η in magnitude (as in (2.2)). Now consider problem (2.4) in this context, which reads as follows:

$$(2.5) \quad \min_{\hat{x}} \left(\max_{\|\delta a\|_2 \leq \eta} \| (a + \delta a) \hat{x} - b \|_2 \right) .$$

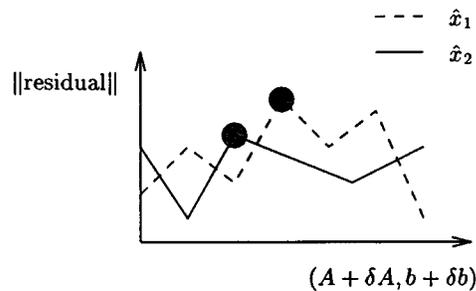


FIG. 2.2. Two illustrative residual-norm curves.

This situation is depicted in Fig. 2.3. The vectors a and b are indicated in thick black lines. The vector a is shown in the horizontal direction and a circle of radius η around its vertex indicates the set of all possible vertices for $a + \delta a$.

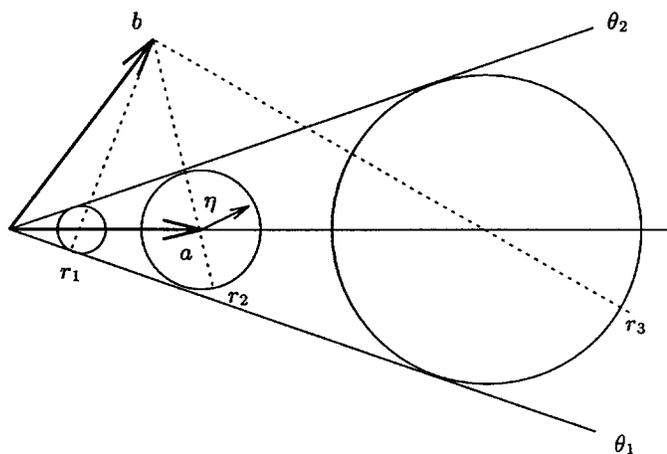


FIG. 2.3. Geometric construction of the solution for a simple example.

For any \hat{x} that we pick, the set $\{(a + \delta a)\hat{x}\}$ describes a disc of center $a\hat{x}$ and radius $\eta\hat{x}$. This is indicated in the figure by the largest rightmost circle, which corresponds to a choice of a positive \hat{x} that is larger than one. The vector in $\{(a + \delta a)\hat{x}\}$ that is furthest away from b is the one obtained by drawing a line from b through the center of the rightmost circle. The intersection of this line with the circle defines a residual vector r_3 whose norm is the largest among all possible residual vectors in the set $\{(a + \delta a)\hat{x}\}$.

Likewise, if we draw a line from b that passes through the vertex of a , it will intersect the circle at a point that defines a residual vector r_2 . This residual will have the largest norm among all residuals that correspond to the particular choice $\hat{x} = 1$.

More generally, any \hat{x} that we pick will determine a circle, and the corresponding largest residual is obtained by finding the furthest point on the circle from b . This is the point where the line that passes through b and the center of the circle intersects the circle on the other side of b .

We need to pick an \hat{x} that minimizes the largest residual. For example, it is clear from the figure that the norm of r_3 is larger than the norm of r_2 . The claim is that

in order to minimize the largest residual we need to proceed as follows: we drop a perpendicular from b to the lower tangent line denoted by θ_1 . This perpendicular intersects the horizontal line in a point where we draw a new circle (the leftmost circle) that is tangent to both θ_1 and θ_2 . This circle corresponds to a choice of \hat{x} such that the furthest point on it from b is the foot of the perpendicular from b to θ_1 . The residual indicated by r_1 corresponds to the desired solution (it has the minimum norm among the largest residuals).

To verify this claim, we refer to Fig. 2.4, where we have only indicated two circles, the circle that leads to a largest residual that is orthogonal to θ_1 and a second circle to its left. For this second leftmost circle, we denote its largest residual by r_4 . We also denote the segment that connects b to the point of tangency of this circle with θ_1 by r . It is clear that r is larger than r_1 since r and r_1 are the sides of a right triangle. It is also clear that r_4 is larger than r by construction. Hence, r_4 is larger than r_1 . A similar argument will show that r_1 is smaller than residuals that result from circles to its right.

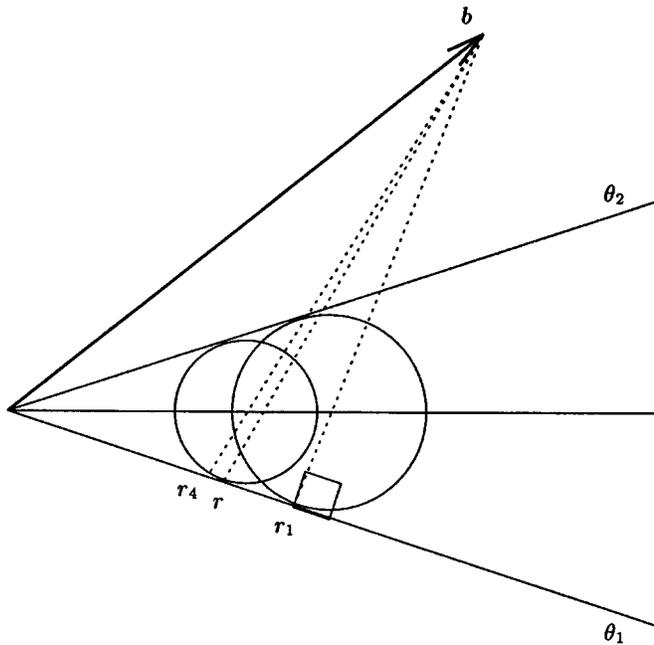


FIG. 2.4. Geometric construction of the solution for a simple example.

The above argument shows that the minimizing solution can be obtained as follows: drop a perpendicular from b to θ_1 . Pick the point where the perpendicular meets the horizontal line and draw a circle that is tangent to both θ_1 and θ_2 . Its radius will be $\eta\hat{x}$, where \hat{x} is the optimal solution. Also, the foot of the perpendicular on θ_1 will be the optimal \hat{b} .

The projection \hat{b} (and consequently the solution \hat{x}) will be nonzero as long as b is not orthogonal to the direction θ_1 . This imposes a condition on η . Indeed, the direction θ_1 will be orthogonal to b only when η is large enough. This requires that the circle centered around a has radius $a^T b$, which is the length of the projection of a onto the unit norm vector b . This is depicted in Fig. 2.5.

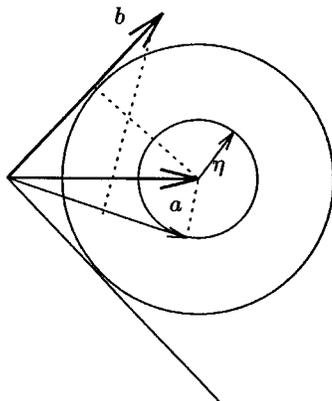


FIG. 2.5. Geometric condition for a nonzero solution.

Hence, the largest value that can be allowed for η in order to have a nonzero solution \hat{x} is

$$\eta < |a^T b|.$$

Indeed, if η were larger than or equal to this value, then the vector in the set $(a + \delta a)$ that would always lead to the maximum residual norm is the one that is orthogonal to b , in which case the solution will be zero again. The same geometric argument will lead to a similar conclusion had we allowed for uncertainties in b as well.

For a nonunity b , the upper bound on η would take the form

$$\eta < \frac{|a^T b|}{\|b\|_2}.$$

We shall see that in the general case a similar bound holds, for nonzero solutions, and is given by

$$\eta < \frac{\|A^T b\|_2}{\|b\|_2}.$$

We now proceed to an algebraic solution of the min-max problem. A final statement of the form of the solution is given in section 3.4.

3. Reducing the min-max problem to a minimization problem. We start by showing how to reduce the min-max problem (2.4) to a standard minimization problem. To begin with, we note that

$$\begin{aligned} \|(A + \delta A)\hat{x} - (b + \delta b)\|_2 &\leq \|A\hat{x} - b\|_2 + \|\delta A\|_2 \cdot \|\hat{x}\|_2 + \|\delta b\|_2, \\ &\leq \|A\hat{x} - b\|_2 + \eta\|\hat{x}\|_2 + \eta_b, \end{aligned}$$

which provides an upper bound for $\|(A + \delta A)\hat{x} - (b + \delta b)\|_2$. But this upper bound is in fact achievable, i.e., there exist $(\delta A, \delta b)$ for which

$$\|(A + \delta A)\hat{x} - (b + \delta b)\|_2 = \|A\hat{x} - b\|_2 + \eta\|\hat{x}\|_2 + \eta_b.$$

To see that this is indeed the case, choose δA as the rank one matrix

$$\delta A^\circ = \frac{(A\hat{x} - b)}{\|A\hat{x} - b\|_2} \frac{\hat{x}^T}{\|\hat{x}\|_2} \eta ,$$

and choose δb as the vector

$$\delta b^\circ = -\frac{(A\hat{x} - b)}{\|A\hat{x} - b\|_2} \eta_b .$$

For these choices of perturbations in A and b , it follows that

$$(A\hat{x} - b) , \quad \delta A^\circ \hat{x} , \quad \text{and} \quad \delta b^\circ ,$$

are collinear vectors that point in the same direction. Hence,

$$\begin{aligned} \|(A + \delta A^\circ) \hat{x} - (b + \delta b^\circ)\|_2 &= \|(A\hat{x} - b) + \delta A^\circ \hat{x} - \delta b^\circ\|_2 , \\ &= \|A\hat{x} - b\|_2 + \|\delta A^\circ \hat{x}\|_2 + \|\delta b^\circ\|_2 , \\ &= \|A\hat{x} - b\|_2 + \eta \|\hat{x}\|_2 + \eta_b , \end{aligned}$$

which is the desired upper bound. We therefore conclude that

$$(3.1) \quad \max_{\|\delta A\|_2 \leq \eta, \|\delta b\|_2 \leq \eta_b} \|(A + \delta A) \hat{x} - (b + \delta b)\|_2 = \|A\hat{x} - b\|_2 + \eta \|\hat{x}\|_2 + \eta_b ,$$

which establishes the following result.

LEMMA 3.1. *The min-max problem (2.4) is equivalent to the following minimization problem. Given $A \in \mathbf{R}^{m \times n}$, with $m \geq n$, $b \in \mathbf{R}^m$, and nonnegative real numbers (η, η_b) , determine, if possible, an \hat{x} that solves*

$$(3.2) \quad \min_{\hat{x}} (\|A\hat{x} - b\|_2 + \eta \|\hat{x}\|_2 + \eta_b) .$$

3.1. Solving the minimization problem. To solve (3.2), we define the cost function

$$\mathcal{L}(\hat{x}) = \|A\hat{x} - b\|_2 + \eta \|\hat{x}\|_2 + \eta_b .$$

It is easy to check that $\mathcal{L}(\hat{x})$ is a convex continuous function in \hat{x} , and hence, any local minimum of $\mathcal{L}(\hat{x})$ is also a global minimum. But at any local minimum of $\mathcal{L}(\hat{x})$, it either holds that $\mathcal{L}(\hat{x})$ is not differentiable or its gradient $\nabla \mathcal{L}(\hat{x})$ is 0. In particular, note that $\mathcal{L}(\hat{x})$ is not differentiable only at $\hat{x} = 0$ and at any \hat{x} that satisfies $A\hat{x} - b = 0$.

We first consider the case in which $\mathcal{L}(\hat{x})$ is differentiable and, hence, the gradient of $\mathcal{L}(\hat{x})$ exists and is given by

$$\begin{aligned} \nabla \mathcal{L}(\hat{x}) &= \frac{1}{\|A\hat{x} - b\|_2} A^T (A\hat{x} - b) + \frac{\eta}{\|\hat{x}\|_2} \hat{x} , \\ &= \frac{1}{\|A\hat{x} - b\|_2} ((A^T A + \alpha I) \hat{x} - A^T b) , \end{aligned}$$

where we have introduced the positive real number

$$(3.3) \quad \alpha = \frac{\eta \|A\hat{x} - b\|_2}{\|\hat{x}\|_2} .$$

By setting $\nabla \mathcal{L}(\hat{x}) = 0$ we obtain that any stationary solution \hat{x} of $\mathcal{L}(\hat{x})$ is given by

$$(3.4) \quad \hat{x} = (A^T A + \alpha I)^{-1} A^T b.$$

We still need to determine the parameter α that corresponds to \hat{x} and which is defined in (3.3).

To solve for α , we introduce the singular value decomposition (SVD) of A ,

$$(3.5) \quad A = U \begin{bmatrix} \Sigma \\ 0 \end{bmatrix} V^T,$$

where $U \in \mathbf{R}^{m \times m}$ and $V \in \mathbf{R}^{n \times n}$ are orthogonal, and $\Sigma = \text{diag}(\sigma_1, \dots, \sigma_n)$ is diagonal, with

$$\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_n \geq 0$$

being the singular values of A . We further partition the vector $U^T b$ into

$$(3.6) \quad \begin{bmatrix} b_1 \\ b_2 \end{bmatrix} = U^T b,$$

where $b_1 \in \mathbf{R}^n$ and $b_2 \in \mathbf{R}^{m-n}$.

In this case, the expression (3.4) for \hat{x} can be rewritten in the equivalent form

$$(3.7) \quad \hat{x} = V(\Sigma^2 + \alpha I)^{-1} \Sigma b_1,$$

and, hence,

$$\|\hat{x}\|_2 = \|\Sigma (\Sigma^2 + \alpha I)^{-1} b_1\|_2.$$

Likewise,

$$\begin{aligned} b - A\hat{x} &= U \left(U^T b - \begin{pmatrix} \Sigma \\ 0 \end{pmatrix} (\Sigma^2 + \alpha I)^{-1} \Sigma b_1 \right), \\ &= U \begin{bmatrix} b_1 - \Sigma^2 (\Sigma^2 + \alpha I)^{-1} b_1 \\ b_2 \end{bmatrix}, \\ &= U \begin{bmatrix} \alpha (\Sigma^2 + \alpha I)^{-1} b_1 \\ b_2 \end{bmatrix}, \end{aligned}$$

which shows that

$$\|b - A\hat{x}\|_2 = \sqrt{\|b_2\|_2^2 + \alpha^2 \|(\Sigma^2 + \alpha I)^{-1} b_1\|_2^2}.$$

Therefore, (3.3) for α reduces to the following nonlinear equation that is only a function of α and the given data (A, b, η) :

$$(3.8) \quad \alpha = \frac{\eta \sqrt{\|b_2\|_2^2 + \alpha^2 \|(\Sigma^2 + \alpha I)^{-1} b_1\|_2^2}}{\|\Sigma (\Sigma^2 + \alpha I)^{-1} b_1\|_2}.$$

Note that only the norm of b_2 , and not b_2 itself, is needed in the above expression.

Remark. We have assumed in the derivation so far that A is full rank. If this were not the case, i.e., if A (and hence Σ) were singular, then (3.8) can be reduced to an equation of the same form but with a nonsingular Σ of smaller dimension. Indeed, if we partition

$$\Sigma = \begin{bmatrix} \hat{\Sigma} & 0 \\ 0 & 0 \end{bmatrix},$$

where $\hat{\Sigma} \in \mathbf{R}^{k \times k}$ is nonsingular, and let $\hat{b}_1 \in \mathbf{R}^k$ be the first k components of b_1 , $\tilde{b}_1 \in \mathbf{R}^{n-k}$ be the last $n - k$ components of b_1 , and let

$$\|\hat{b}_2\|_2^2 = \|b_2\|_2^2 + \|\tilde{b}_1\|_2^2,$$

then (3.8) reduces to

$$(3.9) \quad \alpha = \frac{\eta \sqrt{\|\hat{b}_2\|_2^2 + \alpha^2} \|(\hat{\Sigma}^2 + \alpha I)^{-1} \hat{b}_1\|_2^2}{\|\hat{\Sigma} (\hat{\Sigma}^2 + \alpha I)^{-1} \hat{b}_1\|_2},$$

which is the same form as (3.8). From now on, we assume that A is full rank and, hence, Σ is invertible:

A full rank is a standing assumption in what follows.

3.2. The secular equation. Define the nonlinear function in α ,

$$(3.10) \quad \mathcal{G}(\alpha) = b_1^T (\Sigma^2 - \eta^2 I) (\Sigma^2 + \alpha I)^{-2} b_1 - \frac{\eta^2}{\alpha^2} \|b_2\|_2^2.$$

It is clear that α is a positive solution to (3.8) if, and only if, it is a positive root of $\mathcal{G}(\alpha)$. Following [4], we refer to the equation

$$(3.11) \quad \mathcal{G}(\alpha) = 0$$

as a *secular equation*.

The function $\mathcal{G}(\alpha)$ has several useful properties that will allow us to provide conditions for the existence of a unique positive root α . We start with the following result.

LEMMA 3.2. *The function $\mathcal{G}(\alpha)$ in (3.10) can have at most one positive root. In addition, if $\hat{\alpha} > 0$ is a root of $\mathcal{G}(\alpha)$, then $\hat{\alpha}$ is a simple root and $\mathcal{G}'(\hat{\alpha}) > 0$.*

Proof. We prove the second conclusion first. Partition

$$\begin{bmatrix} \Sigma & 0 \\ 0 & 0 \end{bmatrix} = \begin{bmatrix} \Sigma_1 & 0 \\ 0 & \Sigma_2 \end{bmatrix} \in \mathbf{R}^{(n+1) \times (n+1)},$$

where the diagonal entries of $\Sigma_1 \in \mathbf{R}^{k \times k}$ are those of Σ that are larger than η , and the diagonal entries of $\Sigma_2 \in \mathbf{R}^{(n+1-k) \times (n+1-k)}$ are the remaining diagonal entries of Σ and one 0. It follows that (in terms of the 2-induced norm for the diagonal matrices $(\Sigma_2^2 + \alpha I)$ and $(\Sigma_1^2 + \alpha I)$)

$$(3.12) \quad \|\Sigma_2^2 + \alpha I\|_2 \cdot \|(\Sigma_1^2 + \alpha I)^{-1}\|_2 < 1$$

for all $\alpha > 0$.

Let $u \in \mathbf{R}^k$ be the first k components of $\sqrt{\Sigma^2 - \eta^2 I} \cdot b_1$ and let $v \in \mathbf{R}^{n+1-k}$ be the last $n+1-k$ components of

$$\begin{bmatrix} \sqrt{\eta^2 I - \Sigma^2} & 0 \\ 0 & \eta \end{bmatrix} \begin{bmatrix} b_1 \\ \|b_2\|_2 \end{bmatrix}.$$

It follows that we can rewrite $\mathcal{G}(\alpha)$ as the difference

$$\mathcal{G}(\alpha) = u^T (\Sigma_1^2 + \alpha I)^{-2} u - v^T (\Sigma_2^2 + \alpha I)^{-2} v$$

and, consequently,

$$\mathcal{G}'(\alpha) = -2 \left(u^T (\Sigma_1^2 + \alpha I)^{-3} u - v^T (\Sigma_2^2 + \alpha I)^{-3} v \right).$$

Let $\hat{\alpha} > 0$ be a root of $\mathcal{G}(\alpha)$. This means that

$$u^T (\Sigma_1^2 + \hat{\alpha} I)^{-2} u = v^T (\Sigma_2^2 + \hat{\alpha} I)^{-2} v,$$

which leads to the following sequence of inequalities:

$$\begin{aligned} u^T (\Sigma_1^2 + \hat{\alpha} I)^{-3} u &\leq \|(\Sigma_1^2 + \hat{\alpha} I)^{-1}\|_2 \cdot u^T \cdot (\Sigma_1^2 + \hat{\alpha} I)^{-2} u \\ &= \|(\Sigma_1^2 + \hat{\alpha} I)^{-1}\|_2 \cdot v^T \cdot (\Sigma_2^2 + \hat{\alpha} I)^{-2} v \\ &< \frac{1}{\|(\Sigma_2^2 + \alpha I)\|_2} \cdot v^T \cdot (\Sigma_2^2 + \hat{\alpha} I)^{-2} v \\ &\leq v^T (\Sigma_2^2 + \hat{\alpha} I)^{-3} v. \end{aligned}$$

Combining this relation with the expression for $\mathcal{G}'(\alpha)$, it immediately follows that $\mathcal{G}'(\hat{\alpha}) > 0$. Consequently, $\hat{\alpha}$ must be a simple root of $\mathcal{G}(\alpha)$.

Furthermore, we note that $\mathcal{G}(\alpha)$ is a sum of $n+1$ rational functions in α and hence can have only a finite number of positive roots. In the following we show by contradiction that $\mathcal{G}(\alpha)$ can have no positive roots other than $\hat{\alpha}$. Assume to the contrary that $\hat{\alpha}_1$ is another positive root of $\mathcal{G}(\alpha)$. Without loss of generality, we further assume that $\hat{\alpha} < \hat{\alpha}_1$ and that $\mathcal{G}(\alpha)$ does not have any root within the open interval $(\hat{\alpha}, \hat{\alpha}_1)$. It follows from the above proof that

$$\mathcal{G}'(\hat{\alpha}) > 0 \quad \text{and} \quad \mathcal{G}'(\hat{\alpha}_1) > 0.$$

But this implies that $\mathcal{G}(\alpha) > 0$ for α slightly larger than $\hat{\alpha}$ and $\mathcal{G}(\alpha) < 0$ for α slightly smaller than $\hat{\alpha}_1$, and consequently, $\mathcal{G}(\alpha)$ must have a root in the interval $(\hat{\alpha}, \hat{\alpha}_1)$; a contradiction to our assumptions. So $\mathcal{G}(\alpha)$ can have at most one positive root. \square

Now we provide conditions for $\mathcal{G}(\alpha)$ to have a positive root. (The next result was in fact suggested earlier by the geometric argument of Fig. 2.3.) Note that $A\hat{x}$ can be written as

$$A\hat{x} = U \begin{bmatrix} \Sigma \\ 0 \end{bmatrix} V^T \hat{x}.$$

Therefore solving $A\hat{x} = b$, when possible, is equivalent to solving

$$\begin{bmatrix} \Sigma \\ 0 \end{bmatrix} V^T \hat{x} = U^T b = \begin{bmatrix} b_1 \\ b_2 \end{bmatrix}.$$

This shows that a necessary and sufficient condition for b to belong to the column span of A is $b_2 = 0$.

LEMMA 3.3. *Assume $\eta > 0$ (a standing assumption) and $b_2 \neq 0$, i.e., b does not belong to the column span of A . Then the function $\mathcal{G}(\alpha)$ in (3.10) has a unique positive root if and only if*

$$(3.13) \quad \eta < \frac{\|A^T b\|_2}{\|b\|_2}.$$

Proof. We note that

$$\lim_{\alpha \rightarrow 0^+} (\alpha^2 \mathcal{G}(\alpha)) = -\eta^2 \|b_2\|_2^2 < 0,$$

and that

$$(3.14) \quad \begin{aligned} \lim_{\alpha \rightarrow +\infty} (\alpha^2 \mathcal{G}(\alpha)) &= b_1^T (\Sigma^2 - \eta^2 I) b_1 - \eta^2 \|b_2\|_2^2, \\ &= \|A^T b\|_2^2 - \eta^2 \|b\|_2^2, \\ &= b_1^T \Sigma^2 b_1 - \eta^2 \|b\|_2^2. \end{aligned}$$

First we assume that condition (3.13) holds. It follows then that $\mathcal{G}(\alpha)$ changes sign on the interval $(0, +\infty)$ and therefore has to have a positive root. By Lemma 3.2 this positive root must also be unique.

On the other hand, assume that

$$\eta > \frac{\|A^T b\|_2}{\|b\|_2}.$$

This condition implies, in view of (3.14), that $\mathcal{G}(\alpha) < 0$ for sufficiently large α . We now show by contradiction that $\mathcal{G}(\alpha)$ does not have a positive root. Assume to the contrary that $\hat{\alpha}$ is a positive root of $\mathcal{G}(\alpha)$. It then follows from Lemma 3.2 that $\mathcal{G}(\alpha)$ is positive for α slightly larger than $\hat{\alpha}$ since $\mathcal{G}'(\hat{\alpha}) > 0$, and hence $\mathcal{G}(\alpha)$ must have a root in $(\hat{\alpha}, +\infty)$, which is a contradiction according to Lemma 3.2. Hence $\mathcal{G}(\alpha)$ does not have a positive root in this case.

Finally, we consider the case

$$\eta = \frac{\|A^T b\|_2}{\|b\|_2}.$$

We also show by contradiction that $\mathcal{G}(\alpha)$ does not have a positive root. Assume to the contrary that $\hat{\alpha}$ is a positive root of $\mathcal{G}(\alpha)$. It then follows from Lemma 3.2 that $\hat{\alpha}$ must be a simple root and a continuous function of the coefficients in $\mathcal{G}(\alpha)$. In particular, $\hat{\alpha}$ is a continuous function of η . Now we slightly increase the value of η so that

$$\eta > \frac{\|A^T b\|_2}{\|b\|_2}.$$

By continuity, $\mathcal{G}(\alpha)$ has a positive root for such values of η , but we have just shown that for $\eta > \|A^T b\|_2 / \|b\|_2$ this is not possible. Hence, $\mathcal{G}(\alpha)$ does not have a positive root in this case either. \square

We now consider the case $b_2 = 0$, i.e., b lies in the column span of A . This case arises, for example, when A is a square invertible matrix ($m = n$).

Define

$$\tau_1 = \frac{\|\Sigma^{-1}b_1\|_2}{\|\Sigma^{-2}b_1\|_2} \quad \text{and} \quad \tau_2 = \frac{\|\Sigma b_1\|_2}{\|b_1\|_2} .$$

It follows from $b_2 = 0$ that (cf. (3.13))

$$\tau_2 = \frac{\|A^T b\|_2}{\|b\|_2} .$$

Now note that

$$b_1^T b_1 = b_1^T \Sigma \Sigma^{-1} b_1 .$$

Therefore, by using the Cauchy–Schwarz inequality, we have

$$\|b_1\|_2 \|b_1\|_2 \leq \|\Sigma b_1\|_2 \|\Sigma^{-1} b_1\|_2 ,$$

and we obtain, after applying the Cauchy–Schwarz inequality one more time, that

$$(3.15) \quad \tau_2 = \frac{\|\Sigma b_1\|_2}{\|b_1\|_2} \geq \frac{\|b_1\|_2}{\|\Sigma^{-1} b_1\|_2} \geq \frac{\|\Sigma^{-1} b_1\|_2}{\|\Sigma^{-2} b_1\|_2} = \tau_1 .$$

LEMMA 3.4. *Assume $\eta > 0$ (a standing assumption) and $b_2 = 0$, i.e., b lies in the column span of A . Then the function $\mathcal{G}(\alpha)$ in (3.10) has a positive root if and only if*

$$(3.16) \quad \tau_1 < \eta < \tau_2 .$$

Proof. It is easy to check that

$$\lim_{\alpha \rightarrow 0^+} \mathcal{G}(\alpha) = (\tau_1^2 - \eta^2) b_1^T \Sigma^{-4} b_1 ,$$

and that

$$\lim_{\alpha \rightarrow +\infty} (\alpha^2 \mathcal{G}(\alpha)) = (\tau_2^2 - \eta^2) b_1^T b_1 .$$

If $\eta > \tau_2$, then

$$\lim_{\alpha \rightarrow 0^+} \mathcal{G}(\alpha) < 0 \quad \text{and} \quad \lim_{\alpha \rightarrow +\infty} (\alpha^2 \mathcal{G}(\alpha)) < 0 .$$

Arguments similar to those in the proof of Lemma 3.3 show that $\mathcal{G}(\alpha)$ does not have a positive root. Similarly $\mathcal{G}(\alpha)$ does not have a positive root if $\eta < \tau_1$. Continuity arguments similar to those in the proof of Lemma 3.3 show that $\mathcal{G}(\alpha)$ does not have a positive root if $\eta = \tau_2$ or τ_1 .

However, if $\tau_1 < \eta < \tau_2$, then

$$\lim_{\alpha \rightarrow 0^+} \mathcal{G}(\alpha) < 0 \quad \text{and} \quad \lim_{\alpha \rightarrow +\infty} (\alpha^2 \mathcal{G}(\alpha)) > 0 .$$

So $\mathcal{G}(\alpha)$ must have a positive root. By Lemma 3.2 this positive root is unique. \square

3.3. Finding the global minimum. We now show that whenever $\mathcal{G}(\alpha)$ has a positive root $\hat{\alpha}$, the corresponding vector \hat{x} in (3.4) must be the global minimizer of $\mathcal{L}(\hat{x})$.

LEMMA 3.5. *Let $\hat{\alpha}$ be a positive root of $\mathcal{G}(\alpha)$ and let \hat{x} be defined by (3.4) for $\alpha = \hat{\alpha}$. Then \hat{x} is the global minimum of $\mathcal{L}(\hat{x})$.*

Proof. We first show that

$$\Delta\mathcal{L}(\hat{x}) > 0,$$

where $\Delta\mathcal{L}(\hat{x})$ is the Hessian of \mathcal{L} at \hat{x} . We take the gradient of \mathcal{L} ,

$$\nabla\mathcal{L}(\hat{x}) = \frac{1}{\|A\hat{x} - b\|_2} A^T (A\hat{x} - b) + \frac{\eta}{\|\hat{x}\|_2} \hat{x}.$$

Consequently,

$$\begin{aligned} \Delta\mathcal{L}(\hat{x}) &= \frac{1}{\|A\hat{x} - b\|_2} A^T A - \frac{1}{\|A\hat{x} - b\|_2^3} (A^T A\hat{x} - A^T b) (A^T A\hat{x} - A^T b)^T \\ &\quad + \frac{\eta}{\|\hat{x}\|_2} I - \frac{\eta}{\|\hat{x}\|_2^3} \hat{x}\hat{x}^T. \end{aligned}$$

We now simplify this expression. It follows from (3.4) that

$$(A^T A + \hat{\alpha}I) \hat{x} = A^T b,$$

and, hence,

$$A^T A\hat{x} - A^T b = -\hat{\alpha}\hat{x}.$$

Substituting this relation into the expression for the Hessian matrix $\Delta\mathcal{L}(\hat{x})$, and simplifying the resulting expression using (3.3), we obtain

$$\Delta\mathcal{L}(\hat{x}) = \frac{1}{\|A\hat{x} - b\|_2} \left((A^T A + \hat{\alpha}I) - \frac{\hat{x}\hat{x}^T}{\hat{x}^T \hat{x}} (\hat{\alpha} + \eta^2) \right).$$

Observe that the matrix $(A^T A + \hat{\alpha}I)$ is positive definite since $\hat{\alpha} > 0$. Hence $\Delta\mathcal{L}(\hat{x})$ can have at most one nonpositive eigenvalue. This implies that $\Delta\mathcal{L}(\hat{x})$ is positive definite if and only if $\det(\Delta\mathcal{L}(\hat{x})) > 0$. Indeed,

$$\begin{aligned} \frac{\det(\Delta\mathcal{L}(\hat{x})) \|A\hat{x} - b\|_2^n}{\det(A^T A + \hat{\alpha}I)} &= \det \left(I - \frac{(A^T A + \hat{\alpha}I)^{-1} \hat{x}\hat{x}^T}{\hat{x}^T \hat{x}} (\hat{\alpha} + \eta^2) \right) \\ &= 1 - \frac{\hat{x}^T (A^T A + \hat{\alpha}I)^{-1} \hat{x}}{\hat{x}^T \hat{x}} (\hat{\alpha} + \eta^2) \\ &= \frac{1}{\hat{x}^T \hat{x}} \left(\hat{x}^T \hat{x} - (\hat{\alpha} + \eta^2) \left(\hat{x}^T (A^T A + \hat{\alpha}I)^{-1} \hat{x} \right) \right). \end{aligned}$$

The last expression can be further rewritten, using the SVD of A and (3.8),

$$\begin{aligned} \frac{\det(\Delta\mathcal{L}(\hat{x})) \|A\hat{x} - b\|_2^n}{\det(A^T A + \hat{\alpha}I)} &= \frac{1}{\hat{x}^T \hat{x}} b_1^T \Sigma^2 (\Sigma^2 + \hat{\alpha}I)^{-2} b_1 \\ &\quad - \frac{\hat{\alpha} + \eta^2}{\hat{x}^T \hat{x}} b_1^T \Sigma^2 (\Sigma^2 + \hat{\alpha}I)^{-3} b_1 \end{aligned}$$

$$\begin{aligned}
&= \frac{1}{\hat{x}^T \hat{x}} \frac{\eta^2 \left(\|b_2\|_2^2 + \hat{\alpha}^2 \|(\Sigma^2 + \hat{\alpha}I)^{-1} b_1\|_2^2 \right)}{\hat{\alpha}^2} \\
&\quad - \frac{\hat{\alpha} + \eta^2}{\hat{x}^T \hat{x}} b_1^T \Sigma^2 (\Sigma^2 + \hat{\alpha}I)^{-3} b_1 \\
&= \frac{\hat{\alpha}}{\hat{x}^T \hat{x}} \left(\frac{\eta^2 \|b_2\|_2^2}{\hat{\alpha}^3} + b_1^T (\eta^2 - \Sigma^2) (\Sigma^2 + \hat{\alpha}I)^{-3} b_1 \right).
\end{aligned}$$

Comparing the last expression with the function $\mathcal{G}(\alpha)$ in (3.10), we finally have

$$\frac{\det(\Delta \mathcal{L}(\hat{x})) \|A\hat{x} - b\|_2^q}{\det(A^T A + \hat{\alpha}I)} = \frac{\hat{\alpha}}{2\hat{x}^T \hat{x}} \mathcal{G}'(\hat{\alpha}).$$

By Lemma 3.2, we have that $\mathcal{G}'(\hat{\alpha}) > 0$. Consequently, $\Delta \mathcal{L}(\hat{x})$ must be positive definite, and hence \hat{x} must be a local minimizer of $\mathcal{L}(\hat{x})$. Since $\mathcal{L}(\hat{x})$ is a convex function, this also means that \hat{x} is a global minimizer of $\mathcal{L}(\hat{x})$. \square

We still need to consider the points at which $\mathcal{L}(\hat{x})$ is not differentiable. These include $\hat{x} = 0$ and any solution of $A\hat{x} = b$.

Consider first the case $b_2 \neq 0$. This means that b does not belong to the column span of A and, hence, we only need to check $\hat{x} = 0$. If condition (3.13) holds, then it follows from Lemma 3.3 that $\mathcal{G}(\alpha)$ has a unique positive root $\hat{\alpha}$, and it follows from Lemma 3.5 that

$$\hat{x} = (A^T A + \hat{\alpha}I)^{-1} A^T b$$

is the global minimum. On the other hand, if condition (3.13) does not hold, then it follows from Lemma 3.3 that $\mathcal{G}(\alpha)$ does not have any positive root and hence

$$\hat{x} = 0$$

is the global minimum.

Now consider the case $b_2 = 0$, which means that b lies in the column span of A . In this case $\mathcal{L}(\hat{x})$ is not differentiable at both $\hat{x} = 0$ and $\hat{x} = V\Sigma^{-1}b_1 = A^\dagger b$. If condition (3.16) holds, then it follows from Lemma 3.4 that $\mathcal{G}(\alpha)$ has a unique positive root $\hat{\alpha}$ and it follows from Lemma 3.5 that

$$\hat{x} = (A^T A + \hat{\alpha}I)^{-1} A^T b$$

is the global minimum. On the other hand, if $\eta \leq \tau_1$, then

$$\begin{aligned}
\mathcal{L}(V\Sigma^{-1}b_1) - \mathcal{L}(0) &= \eta \|\Sigma^{-1}b_1\|_2 - \|b_1\|_2, \\
&\leq \|\Sigma^{-1}b_1\|_2 \left(\frac{\|\Sigma^{-1}b_1\|_2}{\|\Sigma^{-2}b_1\|_2} - \frac{\|b_1\|_2}{\|\Sigma^{-1}b_1\|_2} \right), \\
&\leq 0,
\end{aligned}$$

where we have used the Cauchy–Schwarz inequality. It follows that

$$\hat{x} = V\Sigma^{-1}b_1$$

is the global minimum in this case. Similarly, if $\eta \geq \tau_2$, then

$$\hat{x} = 0$$

is the global minimum.

We finally consider the degenerate case $\tau_1 = \tau_2 = \eta$. Under this condition, it follows from (3.15) that

$$\|\Sigma^{-1}b_1\|_2\|\Sigma b_1\|_2 = \|b_1\|_2 \cdot \|b_1\|_2 .$$

Hence,

$$\begin{aligned} \mathcal{L}(V\Sigma^{-1}b_1) - \mathcal{L}(0) &= \eta\|\Sigma^{-1}b_1\|_2 - \|b_1\|_2 \\ &= \frac{\|\Sigma^{-1}b_1\|_2}{\|b_1\|_2} \cdot \|\Sigma^{-1}b_1\|_2 - \|b_1\|_2 = 0 . \end{aligned}$$

This shows that $\mathcal{L}(V\Sigma^{-1}b_1) = \mathcal{L}(0)$. But since $\mathcal{L}(\hat{x})$ is a convex function in \hat{x} , we conclude that for any \hat{x} that is a convex linear combination of 0 and $V\Sigma^{-1}b_1$, say

$$(3.17) \quad \hat{x} = \beta V\Sigma^{-1}b_1, \quad \text{for any } 0 \leq \beta \leq 1 ,$$

we also obtain $\mathcal{L}(\hat{x}) = 0$. Therefore, when $\tau_1 = \tau_2 = \eta$ there are many solutions \hat{x} and these are all scaled multiples of $V\Sigma^{-1}b_1$ as in (3.17).

3.4. Statement of the solution of the min-max problem. We collect, in the form of a theorem, the conclusions of our earlier analysis.

THEOREM 3.6. *Given $A \in \mathbf{R}^{m \times n}$, with $m \geq n$ and A full rank, $b \in \mathbf{R}^m$, and nonnegative real numbers (η, η_b) . The following optimization problem*

$$(3.18) \quad \min_{\hat{x}} \max \{ \| (A + \delta A) \hat{x} - (b + \delta b) \|_2 : \|\delta A\|_2 \leq \eta, \|\delta b\|_2 \leq \eta_b \}$$

always has a solution \hat{x} . The solution(s) can be constructed as follows.

- Introduce the SVD of A ,

$$(3.19) \quad A = U \begin{bmatrix} \Sigma \\ 0 \end{bmatrix} V^T ,$$

where $U \in \mathbf{R}^{m \times m}$ and $V \in \mathbf{R}^{n \times n}$ are orthogonal, and $\Sigma = \text{diag}(\sigma_1, \dots, \sigma_n)$ is diagonal, with

$$\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_n > 0$$

being the singular values of A .

- Partition the vector $U^T b$ into

$$(3.20) \quad \begin{bmatrix} b_1 \\ b_2 \end{bmatrix} = U^T b ,$$

where $b_1 \in \mathbf{R}^n$ and $b_2 \in \mathbf{R}^{m-n}$.

- Introduce the secular function

$$(3.21) \quad \mathcal{G}(\alpha) = b_1^T (\Sigma^2 - \eta^2 I) (\Sigma^2 + \alpha I)^{-2} b_1 - \frac{\eta^2}{\alpha^2} \|b_2\|_2^2 .$$

- Define

$$\tau_1 = \frac{\|\Sigma^{-1}b_1\|_2}{\|\Sigma^{-2}b_1\|_2} \quad \text{and} \quad \tau_2 = \frac{\|A^T b\|_2}{\|b\|_2} .$$

First case: b does not belong to the column span of A .

1. *If $\eta \geq \tau_2$, then the unique solution is $\hat{x} = 0$.*
2. *If $\eta < \tau_2$, then the unique solution is $\hat{x} = (A^T A + \hat{\alpha}I)^{-1} A^T b$, where $\hat{\alpha}$ is the unique positive root of the secular equation $\mathcal{G}(\alpha) = 0$.*

Second case: b belongs to the column span of A .

1. *If $\eta \geq \tau_2$, then the unique solution is $\hat{x} = 0$.*
2. *If $\tau_1 < \eta < \tau_2$, then the unique solution is $\hat{x} = (A^T A + \hat{\alpha}I)^{-1} A^T b$, where $\hat{\alpha}$ is the unique positive root of the secular equation $\mathcal{G}(\alpha) = 0$.*
3. *If $\eta \leq \tau_1$, then the unique solution is $\hat{x} = V\Sigma^{-1}b_1 = A^\dagger b$.*
4. *If $\eta = \tau_1 = \tau_2$, then there are infinitely many solutions that are given by $\hat{x} = \beta V\Sigma^{-1}b_1 = \beta A^\dagger b$, for any $0 \leq \beta \leq 1$.*

The above solution is suitable when the computation of the SVD of A is feasible. For large sparse matrices A , it is better to reformulate the secular equation as follows. Squaring both sides of (3.3) we obtain

$$(3.22) \quad \|(A^T A + \alpha I)^{-1} A^T b\|^2 \alpha^2 = \eta^2 \|A(A^T A + \alpha I)^{-1} A^T b - b\|^2.$$

After some manipulation, we are led to

$$d^T (C + \alpha I)^{-2} d = \frac{\eta^2}{\alpha^2} [b^T b - d^T (C + \alpha I)^{-1} d - \alpha d^T (C + \alpha I)^{-2} d],$$

where we have defined $C = A^T A$ and $d = A^T b$. Therefore, finding α reduces to finding the positive root of

$$(3.23) \quad \begin{aligned} \mathcal{H}(\alpha) &\triangleq d^T (C + \alpha I)^{-2} d \\ &- \frac{\eta^2}{\alpha^2} [b^T b - d^T (C + \alpha I)^{-1} d - \alpha d^T (C + \alpha I)^{-2} d]. \end{aligned}$$

In this form, one can consider techniques similar to those suggested in [5] to find α efficiently.

4. Restricted perturbations. We have so far considered the case in which all the columns of the A matrix are subject to perturbations. It may happen in practice, however, that only selected columns are uncertain, while the remaining columns are known precisely. This situation can be handled by the approach of this paper, as we now clarify.

Given $A \in \mathbf{R}^{m \times n}$, we partition it into block columns,

$$A = \begin{bmatrix} A_1 & A_2 \end{bmatrix},$$

and assume, without loss of generality, that only the columns of A_2 are subject to perturbations while the columns of A_1 are known exactly. We then pose the following min-max problem.

Given $A \in \mathbf{R}^{m \times n}$, with $m \geq n$ and A full rank, $b \in \mathbf{R}^m$, and nonnegative real numbers (η_2, η_b) , determine \hat{x} such that

$$(4.1) \quad \min_{\hat{x}} \max \left\{ \left\| \begin{bmatrix} A_1 & A_2 + \delta A_2 \end{bmatrix} \hat{x} - (b + \delta b) \right\|_2 : \|\delta A_2\|_2 \leq \eta_2, \|\delta b\|_2 \leq \eta_b \right\}.$$

If we partition \hat{x} accordingly with A_1 and A_2 , say

$$\hat{x} = \begin{bmatrix} \hat{x}_1 \\ \hat{x}_2 \end{bmatrix},$$

then we can write

$$\| [A_1 \quad A_2 + \delta A_2] \hat{x} - (b + \delta b) \|_2 = \| (A_2 + \delta A_2) \hat{x}_2 - (b - A_1 \hat{x}_1 + \delta b) \|_2 .$$

Therefore, following the argument at the beginning of section 3, we conclude that the maximum over $(\delta A_2, \delta b)$ is achievable and is equal to

$$\| A_2 \hat{x}_2 - (b - A_1 \hat{x}_1) \|_2 + \eta_2 \| \hat{x}_2 \|_2 + \eta_b .$$

In this way, statement (4.1) reduces to the minimization problem

$$(4.2) \quad \min_{\hat{x}_1, \hat{x}_2} \left(\left\| [A_1 \quad A_2] \begin{bmatrix} \hat{x}_1 \\ \hat{x}_2 \end{bmatrix} - b \right\|_2 + \eta_2 \| \hat{x}_2 \|_2 + \eta_b \right) .$$

This statement can be further reduced to the problem treated in Theorem 3.6 as follows. Introduce the QR decomposition of A , say

$$A = QR = Q \begin{bmatrix} R_{11} & R_{12} \\ 0 & R_{22} \\ 0 & 0 \end{bmatrix} ,$$

where we have partitioned R accordingly with the sizes of A_1 and A_2 . Define

$$\begin{bmatrix} \bar{b}_{1A} \\ \bar{b}_{2A} \\ \bar{b}_2 \end{bmatrix} = Q^T b .$$

Then (4.2) is equivalent to

$$(4.3) \quad \min_{\hat{x}_1, \hat{x}_2} \left(\left\| \begin{bmatrix} R_{11} & R_{12} \\ 0 & R_{22} \\ 0 & 0 \end{bmatrix} \begin{bmatrix} \hat{x}_1 \\ \hat{x}_2 \end{bmatrix} - \begin{bmatrix} \bar{b}_{1A} \\ \bar{b}_{2A} \\ \bar{b}_2 \end{bmatrix} \right\|_2 + \eta_2 \| \hat{x}_2 \|_2 + \eta_b \right) ,$$

which can be further rewritten as

$$(4.4) \quad \min_{\hat{x}_1, \hat{x}_2} \left(\left\| \begin{bmatrix} R_{11} \hat{x}_1 + R_{12} \hat{x}_2 - \bar{b}_{1A} \\ R_{22} \hat{x}_2 - \bar{b}_{2A} \\ \bar{b}_2 \end{bmatrix} \right\|_2 + \eta_2 \| \hat{x}_2 \|_2 + \eta_b \right) .$$

This shows that once the optimal \hat{x}_2 has been determined, the optimal choice for \hat{x}_1 is necessarily the one that annihilates the entry $R_{11} \hat{x}_1 + R_{12} \hat{x}_2 - \bar{b}_{1A}$. That is,

$$(4.5) \quad \hat{x}_1 = R_{11}^{-1} [\bar{b}_{1A} - R_{12} \hat{x}_2] .$$

The optimal \hat{x}_2 is the solution of

$$(4.6) \quad \min_{\hat{x}_2} \left(\left\| \begin{bmatrix} R_{22} \\ 0 \end{bmatrix} \hat{x}_2 - \begin{bmatrix} \bar{b}_{2A} \\ \bar{b}_2 \end{bmatrix} \right\|_2 + \eta_2 \| \hat{x}_2 \|_2 + \eta_b \right) .$$

This optimization is of the same form as the problem stated earlier in Lemma 3.1 with \hat{x} replaced by \hat{x}_2 , η replaced by η_2 , A replaced by $\begin{bmatrix} R_{22} \\ 0 \end{bmatrix}$, and b replaced by $\begin{bmatrix} \bar{b}_{2A} \\ \bar{b}_2 \end{bmatrix}$.

Therefore, the optimal \hat{x}_2 can be obtained by applying the result of Theorem 3.6. Once \hat{x}_2 has been determined, the corresponding \hat{x}_1 follows from (4.5).

5. Conclusion. In this paper we have proposed a new formulation for parameter estimation in the presence of data uncertainties. The problem incorporates a priori bounds on the size of the perturbations and admits a nice geometric interpretation. It also has a closed form solution that is obtained by solving a regularized least-squares problem with a regression parameter that is the unique positive root of a secular equation.

Several other interesting issues remain to be addressed. Among these, we state the following:

1. A study of the statistical properties of the min-max solution is valuable for a better understanding of its performance in stochastic settings.
2. The numerical properties of the algorithm proposed in this paper need also be addressed.
3. Extensions of the algorithm to deal with perturbations in submatrices of A are of interest and will be studied elsewhere.

We can also extend the approach of this paper to other variations that include uncertainties in a weighting matrix, multiplicatives uncertainties, etc. (see, e.g., [15]).

REFERENCES

- [1] L. E. GHAOUI AND H. LEBRET, *Robust solutions to least-squares problems with uncertain data*, SIAM J. Matrix Anal. Appl., 18 (1997), pp. 1035–1064.
- [2] G. H. GOLUB, *Some modified matrix eigenvalue problems*, SIAM Rev., 15 (1973), pp. 318–344.
- [3] G. H. GOLUB AND C. F. VAN LOAN, *An analysis of the total least squares problem*, SIAM J. Numer. Anal., 17 (1980), pp. 883–893.
- [4] G. H. GOLUB AND C. F. VAN LOAN, *Matrix Computations*, 3rd ed., The Johns Hopkins University Press, Baltimore, MD, 1997.
- [5] G. H. GOLUB AND U. VON MATT, *Generalized Cross-Validation for Large Scale Problems*, Tech. report SCCM-96-06, Computer Science Department, Stanford University, Stanford, CA, 1996.
- [6] B. HASSIBI, A. H. SAYED, AND T. KAILATH, *Linear estimation in Krein spaces - Part I: Theory*, IEEE Trans. Automat. Control, 41 (1996), pp. 18–33.
- [7] S. V. HUFFEL AND J. VANDEWALLE, *The Total Least Squares Problem: Computational Aspects and Analysis*, SIAM, Philadelphia, PA, 1991.
- [8] S. M. KAY, *Fundamentals of Statistical Signal Processing*, Prentice-Hall, Englewood Cliffs, NJ, 1993.
- [9] P. KHARGONEKAR AND K. M. NAGPAL, *Filtering and smoothing in an H^∞ - setting*, IEEE Trans. Automat. Control, AC-36 (1991), pp. 151–166.
- [10] C. L. LAWSON AND R. J. HANSON, *Solving Least-Squares Problems*, SIAM, Philadelphia, PA, 1995. Revised republication of work first published by Prentice-Hall, Englewood Cliffs, NJ, 1974.
- [11] L. LJUNG AND T. SÖDERSTRÖM, *Theory and Practice of Recursive Identification*, MIT Press, Cambridge, MA, 1983.
- [12] A. H. SAYED, B. HASSIBI, AND T. KAILATH, *Fundamental inertia conditions for the minimization of quadratic forms in indefinite metric spaces*, in Recent Developments in Operator Theory and Its Applications, Oper. Theory Adv. Appl. 87, Birkhäuser, Basel, 1996, pp. 309–347.
- [13] L. L. SCHARF, *Statistical Signal Processing: Detection, Estimation and Time Series Analysis*, Addison-Wesley, Reading, MA, 1991.
- [14] U. SHAKED AND Y. THEODOR, *H^∞ -optimal estimation: A tutorial*, in Proc. IEEE Conference on Decision and Control, Tucson, AZ, December 1992, IEEE Computer Society Press, Los Alamitos, CA, pp. 2278–2286.
- [15] S. CHANDRASEKARAN, G. GOLUB, M. GU, AND A. H. SAYED, *Parameter estimation in the presence of bounded modeling errors*, IEEE Signal Processing Letters, 4 (1997), pp. 195–197.