Models for Robust Estimation and Identification

S. Chandrasekaran[1] and K. E. Schubert[2]

[1] Department of Electrical and Computer Engineering
University of California
Santa Barbara, CA 93106
[2] Department of Mathematics and Computer Science
University of Redlands
Redlands, CA 92374

Abstract

In this paper, we will investigate estimation and identification theories with the goal of determining some new methods of adding robustness. We consider uncertain estimation problems, namely ones in which the uncertainty multiplies the quantities to be estimated. Mathematically the problem can be stated as, for system matrices and data matrices that lie in the sets $(A + \delta A)$ and $(b + \delta b)$ respectively, find the value of x that minimizes the cost $\|(A + \delta A)x - (b + \delta b)\|$. We will examine how the proposed techniques compare with currently used methods such as Least Squares (LS), Total Least Squares (TLS), and Tikhonov Regularization (TR). Several results are presented and some future directions are suggested.

1 Introduction

Every system in the real world is uncertain to one degree or another, and thus everyone who does estimation or identification must consider the assumptions made about the system and resulting problem. Consider, for example the simple system described by

$$Ax = b,$$

with

$$\begin{aligned} A &= \begin{bmatrix} 0.11765 & 0.12909 \\ -0.24957 & -0.26919 \end{bmatrix} \\ b &= \begin{bmatrix} -0.074888 \\ 0.154728 \end{bmatrix}. \end{aligned}$$

[0]E-mail addresses: shiv@ece.ucsb.edu and schubert@uor.edu. The authors were partially supported by NSF grant CCR-9734290.

For this exact system the solution is given by

$$x \quad = \quad \begin{bmatrix} 0.34 \\ -0.89 \end{bmatrix}.$$

This is a nice system with reasonable condition number, but if we round A to two decimal places,

$$A \quad = \quad \begin{bmatrix} 0.12 & 0.13 \\ -0.25 & -0.27 \end{bmatrix},$$

we find that the new solution is

$$x \quad = \quad \begin{bmatrix} 1.0505 \\ -1.5457 \end{bmatrix}.$$

The best thing we can say about this is that the signs of the solution are correct. This illustrates that even innocent looking systems can exhibit bad behavior in normal situations. What can be done? Consider the general form of the regularized solution,

$$x(\psi) = \left(A^T A + \psi I\right)^{-1} A^T b, \tag{1}$$

with $\psi = 10^{-7}$. This yields a solution of

$$x\left(10^{-7}\right) \quad = \quad \begin{bmatrix} 0.21515 \\ -0.77273 \end{bmatrix}.$$

This is much better, but can we automate the selection of the regularization parameter? We would like to examine the one parameter family given in equation 1 to find the one closest to the true system, but that requires knowing the answer apriori. We must look for another standard. Note that frequently the exact solution is not as important as the residual so we need to take this into account. We will consider several methods for selecting the regulation parameter in this paper. We will then end with a numerical example of performance.

2 Minmax

The minmax problem was described and solved in [2, 9]. The underlying assumption is that the true system is the worst possible one in the entire set of all possible systems specified by the uncertainties in the problem. The assumptions are thus clearly pessimistic. Let's look at the results of the model. The basic problem can be stated as

$$\min_x \max_{\substack{\|\delta A\| \leq \eta \\ \|\delta b\| \leq \eta_b}} \left(\|(A + \delta A)x - (b + \delta b)\|\right).$$

After performing the maximization the resulting problem is

$$\min_x \left(\|Ax - b\| + \eta\|x\| + \eta_b\right).$$

The problem is convex so the solution is guaranteed, but involves several cases. The basic structure of the solution can be stated as

$$\begin{aligned} x_{minmax} &= (A^TA+\alpha I)^{-1}A^Tb \\ &\alpha = \frac{\eta\|Ax-b\|}{\|x\|}. \end{aligned}$$

The majority of work for this problem revolves around finding the value of α. One thing we can assert about α though is that it is non-negative and in general it will be strictly positive. We note that as a result of this we can state that,

$$\frac{\sigma_n^2}{\sigma_n^2+\alpha}\|x_{LS}\| \leq \|x_{minmax}\| \leq \frac{\sigma_1^2}{\sigma_1^2+\alpha}\|x_{LS}\|.$$

The result follows by taking norms of the expression for the minmax solution and doing some algebra. The point of this is that the minmax returns an estimate that has less signal strength (smaller norm) than the well known LS estimator. The assumptions and results are thus clearly pessimistic.

3 Multi-Column Minmax

A special case of this problem was solved in [2], where one block column of the matrix A is assumed known exactly. Essentially, this problem also assumes that the real system is worse, but now we can impose column bounds, which can reflect relative certainty of different parts in the system. The problem is still pessimistic, but contains several twists on the non-partitioned problem. The general problem, with A partitioned into n block columns, can be stated as

$$\min_{x_i} \max_{\substack{\|\delta A_i\| \leq \eta_i \\ \|\delta b\| \leq \eta_b}} \left(\|\sum_{i=1}^{n}(A_i+\delta A_i)x_i-(b+\delta b)\| \right).$$

Using techniques similar to those in [2], we can simplify the problem to

$$\min_{x_i} \left(\|Ax-b\| + \sum_{i=1}^{n}\eta_i\|x_i\| + \eta_b \right).$$

This problem is a convex sum of euclidean norms. A large body of literature exists for solving the sum of Euclidean norms problem. The problem dates back to Fermat, who posed a special case. Various methods have been proposed which range from a sequence of linear least squares problems [15, 10, 6, 7, 14] to successive over-relaxation [13] to hyperbolic approximation procedure [8] to subgradients [11, 5]. All of these have at best linear convergence so we recommend the quadratically convergent method proposed by Michael Overton in [12]. Overton's method uses an active set and considers the projected objective function which is locally continuously differentiable. Note that Overton's method is similar to [1].

All of the methods mentioned do not take advantage of the reduction in size that can be obtained by using a secular equation as was done in the simpler non-partitioned case. We note that the basic solution when there are p partitions can be written as

$$\begin{aligned} x_i &= \frac{1}{\psi_i} A_i^T \left(I + \sum_{i=1}^{p} \frac{1}{\psi_i} A_i A_i^T \right)^{-1} b \\ &= \frac{1}{\psi_i} A_i^T \left(I + A \Psi^{-1} A^T \right)^{-1} b \end{aligned}$$

with

$$\begin{aligned} \Psi &= \begin{bmatrix} \psi_1 I & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & \psi_p I \end{bmatrix} \\ \psi_i &= \frac{\eta_p \left\| A\hat{x} - b \right\|}{\left\| \hat{x}_p \right\|} \end{aligned}$$

and

$$x_\Psi = \begin{bmatrix} x_1^T & \cdots & x_p^T \end{bmatrix}^T .$$

We note that we can also write a secular equation for each partition, which can be solved to find the values of ψ_i. The secular equations, G_i, are given by

$$G_i(\psi) = b^T F N_i F b \tag{2}$$

with

$$\begin{aligned} F &= \left(I + \sum_{i=1}^{p} \frac{1}{\psi_i} A_i A_i^T \right)^{-1}, \\ N_i &= \left(A_i A_i^T - \eta_i^2 I \right). \end{aligned}$$

While this is a smaller problem than the original, and in most cases can be solved rapidly, we do not have a proof of quadratic convergence of the root finder for the secular equations. It is thus not guaranteed to be faster, but offers a potential savings.

On a different front, we would think that the size of the multi-column partitioned min-max solution, x_Ψ should be smaller than the least squares solution, x_{LS}, since both have the same numerator and the denominator of x_Ψ is larger. This seems reasonable particularly given that this was true in the non-partitioned case, and in some sense the partitioned case reflects a problem that is more known and thus less uncertain. This is not always the case though. To demonstrate this we consider a simple problem. Consider the following A and

b matrices with each column of A a separate partition,

$$A = \begin{bmatrix} 1 & 0 & 0.1 \\ 1 & -1 & 1 \\ 0 & 0 & 0.1 \\ 0 & 0 & 0 \end{bmatrix} \quad b = \begin{bmatrix} 1 \\ 0 \\ 0 \\ 10 \end{bmatrix}.$$

We note that the least squares solution is given by

$$x_{LS} = [1 \quad 1 \quad 0]^T.$$

Now consider the case when $\eta_1 = 2$,$\eta_2 = 0$, and $\eta_3 = 0$. The solution, x_Ψ is given by

$$x_\Psi = [0 \quad 5 \quad 5]^T.$$

It is trivial to see that $\|x_{LS}\| < \|x_\Psi\|$, and thus the idea is disproved. The question remains then as to what we can say about the size of x_Ψ and thus where it lies. The following is not tight in its bound but it does provide a good starting comparison to the non-partitioned case which always has a smaller norm than the least squares,

$$\|x_\Psi\| \leq \frac{\sigma_1^2}{\sigma_n^2}\|x_{LS}\|.$$

Additionally, in the non-partitioned problem we have the simple relation that the solution x is non-zero if and only if $\|A^T b\| > \eta\|b\|$. This is not true for the partitioned case. This is easily seen by considering the following:

$$A_1 = \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix} \quad A_2 = \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix} \quad b = \begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix}.$$

It is readily apparent that $A_2^T b = 0$ and thus from the original problem we should have $x_2 = 0$ for all η_2. Now consider $\eta_1 = \eta_2 = \frac{1}{4}$ and we find that

$$\begin{aligned} x_1 &= 1 - \frac{2}{\sqrt{11}} \\ x_2 &= \frac{3}{\sqrt{11}} - 1. \end{aligned}$$

Our intuition from the non-partitioned case suggests that $x_2 = 0$, but this does not hold because of the column interaction. The partitioned and non-partitioned cases are fundamentally different.

So what do we have? First, we have a quadratically convergent method for finding the solution, as provided by Overton's method. Second, we have a region which contains the solution. Third, we have the form of solution and a secular equation for each partition. Fourth, we can see that the solution is surprisingly different from the non-partitioned case, and so applying results from one case to the other is inherently dangerous.

4 Bounded Errors-in-Variables

This problem has a degenerate (multiple solution) case and a non-degenerate case. The non-degenerate case was solved in [3], while the degenerate case was solved in [4]. The problem assumes the underlying problem is the best possible, similar to TLS but with a bound on how far the A and b matrices can be projected. This is the first optimistic problem we consider, and it can be stated as

$$\min_{x} \min_{\substack{\|\delta A\| \leq \eta \\ \|\delta b\| \leq \eta_b}} \left(\|(A + \delta A)x - (b + \delta b)\| \right).$$

Additionally this problem is not convex. In [3], it was proven that the necessary and sufficient conditions for non-degeneracy are

1. $\eta < \sigma_n$,
2. $b^T(I - A(A^T A - \eta^2 I)^{-1}b > 0$.

The problem can be reduced to

$$\min_{x}(\|Ax - b\| - \eta\|x\|).$$

Note that in the degenerate case the additional constraint of selecting the x with minimum norm is imposed in [4] to get a unique solution. The general form of the solution for the non-degenerate case is

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

with

$$\alpha = \frac{\eta\|Ax - b\|}{\|x\|}.$$

The value of α is determined by a secular equation. Note that the solution in the non-degenerate case always does de-regulation. On the other hand the solution in the degenerate case is

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

with

$$\max(-\sigma_n^2, -\eta^2) \leq \alpha \leq \eta\sigma_1.$$

The particular value of α is given by a secular equation. Here we can see that if $\alpha > 0$ then the degenerate case will do regulation, so the degenerate case can either de-regularize (be optimistic) or regularize (be pessimistic). It is also interesting to note that the degenerate bounded errors-in-variables and the minmix models can sometimes give the same solution. In this case the solution has the best features of both methods.

5 Backward Error

The final problem we will consider is the backward error model. This model contains both optimistic and pessimistic assumptions, and is non-convex. This problem is taken up in a paper to be submitted shortly for publication. The problem is given by the expression

$$\min_x \max_{\|\delta A\| \leq \eta} \frac{\|(A+\delta A)x - b\|}{\|A\|\|x\| + \|b\|}.$$

The maximization can be performed without difficulty to obtain

$$\min_x \frac{\|Ax - b\| + \eta\|x\|}{\|A\|\|x\| + \|b\|}.$$

Due to the difficulty of the problem, we pose instead an intermediate problem that demonstrates some interesting qualities of the original.

$$\min_x \frac{\|Ax - b\| + \eta\|x\|}{\|A\|\|x\|}.$$

Note that the solution to this problem is identical to

$$\min_x \frac{\|Ax - b\|}{\|A\|\|x\|}.$$

The solution is found by taking the derivative and setting equal to zero. The resulting solution form is denoted SBE for simplified backward error and is given by

$$x_{SBE} = \left(A^T A - \gamma_{SBE} I\right)^{-1} A^T b$$

with

$$\gamma_{SBE} = \frac{\|Ax_{SBE} - b\|^2}{\|x_{SBE}\|^2}.$$

The particular value of γ_{SBE} is determined by the root of a secular equation in the interval, $0 \leq \gamma_{SBE} \leq \sigma_n^2$. We can see that these problems de-regularize, and so contain optimistic assumptions. We can even tighten up the interval to show that

$$\sigma_{n+1}^2 \leq \gamma \leq \sigma_n^2,$$

where σ_{n+1} is the TLS parameter. Thus the simpler problem is more optimistic than TLS! One repercussion of the lack of the norm of b in the denominator of the cost is that it is possible for one element of the solution of the simple problem to become infinite in a particular case. The full backward error problem is thus more desirable. Generally, the smaller the regression parameter, the better the result. In most cases the full backward error produces the smallest regression parameter, and thus tends to give the best solution.

6 Numerical Example

We have discussed several different problem formulations that can be used in estimation. We now want to get a feel for how these problems operate on a simple example. Consider for example a simple one dimensional "skyline" image that has been blurred. A "skyline" image is a one dimensional image that looks like a city skyline when graphed, and thus is the most basic image processing example. "Skyline" images involve sharp corners, and it is of key importance to accurately locate these corner transitions. Blurring occurs often in images. For example atmospheric conditions, dust, or imperfections in the optics can cause a blurred image. Blurring is usually modeled as a gaussian blur, which is a great smoothing filter. The gaussian blur causes greater distortion on the corners, which is exactly where we do not want it. The component of a gaussian blur with standard deviation, σ, in position, (i,j), is given by

$$G_{i,j} = e^{-\left(\frac{i-j}{\hat{\sigma}}\right)^2}.$$

If we go on the presumption that we do not know the exact blur that was applied (the standard deviation, $\hat{\sigma}$ unknown) we cannot expect to get the exact image back. While we realize that we will not be able to perfectly extract the original system, we want to see if we can get a little more information than we have now. We "know" the blur is small compared to the information so we are confident that we should be able to get something. The least squares solution fails completely, yielding a result that is about three orders of magnitude off, see Figure 1. We notice that the TLS solution is better than the LS solution, but still not acceptable. The ridge regression solution works well due to its increased robustness. The minmax performs well given its robustness. We see that the bounded errors-in-variables solution also exhibits robustness as this is one of the cases where the problem is degenerate and it can either regularize or de-regularize. In this case the solution is regularized due to the relatively large uncertainty. Most interestingly note that the backward error solution performs the best of all. It does an excellent job of finding the corners without sacrificing the edges.

References

[1] P. Calamai and A. Conn, *A stable algorithm for solving the multifacility location problem involving euclidean distances*, SIAM Journal on Scientific and Statistical Computing, 1 (1980), pp. 512–526. Another quadratically convergent method for the sum of Euclidean norms.

[2] S. Chandrasekaran, G. H. Golub, M. Gu, and A. H. Sayed, *Parameter estimation in the presence of bounded data uncertainties*, SIMAX, 19 (1998), pp. 235–252.

[3] ——, *An efficient algorithm for a bounded errors-in-variables model*, SIMAX, 20 (1999), pp. 839–859.

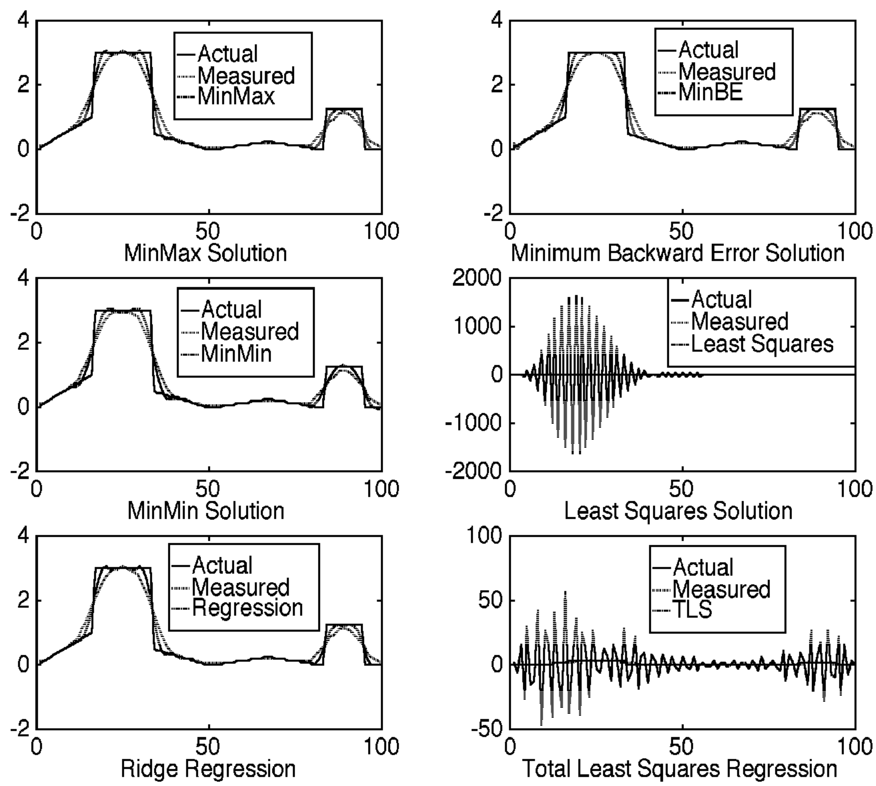

Figure 1: Skyline Problem

[4] S. Chandrasekaran, M. Gu, A. H. Sayed, and K. E. Schubert, *The degenerate bounded errors-in-variables model*, SIMAX, 23 (2001), pp. 138–166.

[5] F. Cordellier and J. Giorot, *On the fermat-weber problem with convex cost functionals*, Mathematical Programming, 14 (1978), pp. 295–311. Convex Programming method of solving multi-column problem.

[6] U. Eckhardt, *On an optimization problem related to minimal surfaces with obstacles*, in Optimization and optimal control, R. Bulirsch, W. Oettli, and J. Stoer, eds., Springer-Verlag, 1975. More work on solving sum of euclidean norms problem by series of linear least squares.

[7] ——, *Weber's problem and weiszfeld's algorithm in general spaces*, Mathematical Programming, 18 (1980), pp. 186–196. Solve multi-column problem by a series of linear least squares problems.

[8] J. Eyster, J. White, and W. Wierwille, *On solving multifacility location problems using a hyperboloid approximation procedure*, AIIE Transactions, 5 (1973), pp. 1–6. Solve multi-column problem by hyperboloid approximation procedure.

[9] L. E. Ghaoui. and H. Lebret, *Robust solutions to least-squares problems with uncertain data*, SIMAX, 18 (1997), pp. 1035–1064.

[10] H. Kuhn, *A note on fermat's problem*, Mathematical Programming, 4 (1973), pp. 98–107. Advocates Weiszfeld's series of linear least squares problems approach to sum of Euclidean norms problem.

[11] R. Love, *Locating facilities in three-dimensional space by convex programming*, Naval Research Logistics Quarterly, 16 (1969), pp. 503–516. Uses convex programming to solve sum of Euclidean norms problem.

[12] M. L. Overton, *A quadratically convergent method for minimizing a sum of euclidean norms*, Mathematical Programming, 27 (1983), pp. 34–63.

[13] S. Schechter, *Minimization of a convex function by relaxation*, in Integer and nonlinear programming, J. Abadie, ed., North-Holland, 1970, pp. 177–190. More work on solving sum of euclidean norms problem by series of linear least squares.

[14] H. Voss and U. Eckhardt, *Linear convergence of generalized weiszfeld's method*, Computing, 25 (1980), pp. 243–251. Proves some convergence properties for Weiszfeld's method.

[15] E. Weiszfeld, *Sur le point par lequel la somme des distances de n points donns est minimum*, Tohoku Mathematics Journal, 43 (1937), pp. 355–386. Not in English, but the main reference to modern solutions to sum of Euclidean norms problem.