MODELS FOR ROBUST ESTIMATION AND IDENTIFICATION

S. Chandrasekaran[1]
Department of
Electrical and Computer Engineering
University of California
Santa Barbara, CA 93106
shiv@ece.ucsb.edu

K. E. Schubert[1]
Department of
Computer Science
California State University
San Bernardino, CA 92407
schubert@csci.csusb.edu

ABSTRACT: *In this paper, estimation and identification theories will be examined with the goal of determining some new methods of adding robustness. The focus will be upon 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 $\|(A+\delta A)x - (b+\delta b)\|$. The proposed techniques are compared 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.*

Keywords: regression, uncertainty, robustness, errors-in-variables, total least squares

1. INTRODUCTION

Consider the set of linear equations, $Ax = b$, where $A \in \mathbb{R}^{m\times n}$ and $b \in \mathbb{R}^m$ are given. The goal is to calculate the value of $x \in \mathbb{R}^n$. If the equation is exact and A is not singular, the solution can be readily found by a variety of techniques, such as taking the QR factorization of A.

$$\begin{aligned} Ax &= b \\ QRx &= b \\ Rx &= Q^T b \end{aligned}$$

The last equation can be solved for x by back-substitution, since R is upper triangular. Given errors in modeling, estimation, and numeric representation the equality rarely holds. The least squares (LS) technique directly uses techniques like the QR factorization, by considering all the errors to be present in b. A more realistic appraisal of the system, considers errors in both A and b. This paper will assume that the "true" system is $(A+\delta A)$ and $(b+\delta b)$, with some condition used to specify the uncertainties, δA and δb. The general description of a system with uncertainty can be expressed as

$$\min_x \|(A+\delta A)x - (b+\delta b)\| \tag{1}$$

given the conditions on δA and δ.

It might seem that this is only for special cases, however, 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

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

For this exact system the solution is given by

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

This is a nice system with a reasonable condition

[1]The authors were partially supported by NSF grant CCR-9734290.

number, but if A is rounded to two decimal places,

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

the new solution is

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

The best thing that can be said 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,

$$x(\psi) = \left(A^T A + \psi I\right)^{-1} A^T b, \tag{2}$$

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

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

This is much better, but can the selection of the regularization parameter be automated? It would be ideal to examine the one parameter family given in Eq. 2 to find the member closest to the true system, but that requires knowing the answer a priori. The only way to handle the numeric difficulty without a priori information is to use the formulation of Eq. 1 to account for the uncertainty in the overall problem.

A well known technique that addresses this problem is total least squares (TLS), [15, 13, 12, 10, 9]. TLS presumes that the original system is consistent and thus it defines the uncertainty by $(b + \delta b) \in \mathcal{R}(A + \delta A)$. TLS selects $\psi = \sigma_{n+1}$, where σ_{n+1} is the smallest singular value of $[A \quad b]$. This works well unless the system is ill-conditioned, which unfortunately is common. Alterations to the basic problem such as truncating the singular values, [10], or integrating Tikhonov estimation with TLS, [12], have been tried but require ad-hoc methodologies.

Another standard must be found. Several methods for selecting the regulation parameter will be examined in this paper. Finally, two numerical examples of performance will be considered.

2. MIN MAX

The min max problem was described and solved in [2, 11]. 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 model can be stated as

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

After performing the maximization an equivalent formulation is

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

The equation is convex so the solution is guaranteed, but involves several cases. The general form of the solution can be stated as

$$x_{minmax} = (A^T A + \psi I)^{-1} A^T b \tag{3}$$

$$\psi = \frac{\eta\|Ax - b\|}{\|x\|}. \tag{4}$$

To find ψ a secular equation[1] is formed by rewriting Eq. 4 using the singular value decomposition (SVD) of A and the expression for x_{minmax} in Eq. 4. The secular equation is thus

$$g(\psi) = b_1^T(\Sigma^2 - \eta^2 I)(\Sigma^2 + \psi I)^{-2} b_1 - \frac{\eta^2}{\psi^2}\|b_2\|^2$$

$$\begin{bmatrix} U_1 & U_2 \end{bmatrix}^T b = \begin{bmatrix} b_1 \\ b_2 \end{bmatrix}$$

$$A = \begin{bmatrix} U_1 & U_2 \end{bmatrix} \begin{bmatrix} \Sigma \\ 0 \end{bmatrix} V^T.$$

The resulting equation can be easily solved using a root finder, such as Newton's method or bisection. One thing we can assert about ψ 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,

$$\frac{\sigma_n^2}{\sigma_n^2 + \psi}\|x_{LS}\| \leqslant \|x_{minmax}\| \leqslant \frac{\sigma_1^2}{\sigma_1^2 + \psi}\|x_{LS}\|.$$

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

[1] This follows the naming scheme used by Golub and Van Loan in [14].

3. MULTI-COLUMN MIN MAX

In the multiple (block) column partitioning case A and δA are considered to be partitioned into block columns (A_j and δA_j), and the norm of each partition of the error, δA_j, is assigned a bound. A special case of this problem was solved in [2], where one block column of the matrix A is assumed to be known exactly ($\delta A_2 = 0$). The problem can be thought of as an order updating problem if partitioned into two blocks. For multiple block columns, this method is useful in tracking multiple targets on radar arrays, or working with inverse problems such as those in seismology. The problem, for p block partitions, is given as

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

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

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

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 [21, 16, 6, 7, 20] to successive over-relaxation [19] to hyperbolic approximation procedure [8] to subgradients [17, 5]. All of these have at best linear convergence so we recommend the quadratically convergent method proposed by Michael Overton in [18]. 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

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

and

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

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

$$G_i(\psi) = b^T F N_i F b \tag{5}$$

with

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

This is a smaller problem than the original, and in most cases can be solved rapidly, using a multi-dimensional root finder.

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,

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

We note that the least squares solution is given by

$$x_{LS} = \begin{bmatrix} 1 & 1 & 0 \end{bmatrix}^T.$$

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

$$x_\Psi = \begin{bmatrix} 0 & 5 & 5 \end{bmatrix}^T.$$

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,

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

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 (thus a reduced size technique to find the solution). 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. MIN MIN

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. The method can be stated as

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

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

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

with

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

Note that the solution in the non-degenerate case always does de-regulation. On the other hand the solution in the degenerate case is

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

with

$$\max(-\sigma_n^2, -\eta^2) \leqslant \psi \leqslant \eta\sigma_1.$$

The particular value of ψ is given by a secular equation. Here we can see that if $\psi > 0$ then the degenerate case will do regulation, so the degenerate case can either de-regularize or regularize. It is also interesting to note that the degenerate min min and the min mix 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 nonconvex. This problem is taken up in a paper to be submitted shortly for publication. The problem is given by the expression

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

The maximization can be performed 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.

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

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 it equal to zero. The resulting solution form is denoted SBE for simplified backward error and is given by

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

with

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

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

$$\sigma_{n+1}^2 \leqslant \psi \leqslant \sigma_n^2,$$

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 EXAMPLES

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 two examples from image processing. Blurring occurs often in images. For example atmospheric conditions, dust, or imperfections in the optics can cause a blurred image. Blurring is usually modelled 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

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

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.

6.1. First Example

Consider 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. The LS solution captures general trends, but still not acceptable, see Figure 1. The Tikhonov solution works well due to its increased robustness. Observe that the min min solution exhibits robustness as this is one of the cases where the problem

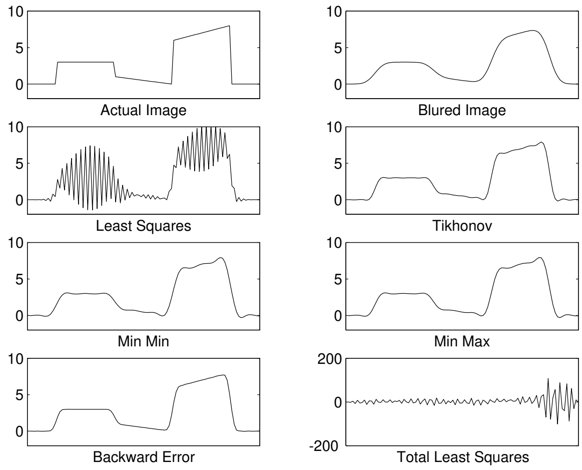

Figure 1: Skyline Problem

is degenerate and it can either regularize or deregularize. In this case the solution is regularized due to the relatively large uncertainty. The min max performs well due to its robustness. 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. Finally, the TLS solution fails completely, yielding a result that is about two orders of magnitude off.

6.2. Second Example

The second example is a simple two-dimensional image processing application. A small picture with the grey-scale words, 'HELLO WORLD' of early programming fame, has been blurred. The image is 20x35 and the blur is done by a Gaussian blur matrix of size 20. The blur is not so strong that some of the features cannot be seen, and in particular one can see that there is writing but the specifics are hard to make out.

A key aspect of all of the regression techniques is selection of the regression parameter. In the suggested techniques, this is done semi-automatically. Semi-automatically because the error bound on the matrix must still be supplied, and it is not guaranteed to be known accurately. This becomes critical as the selection of the regression parameter is mostly influenced by this error bound. Select an error that is too large and data losses can result, select one too small and

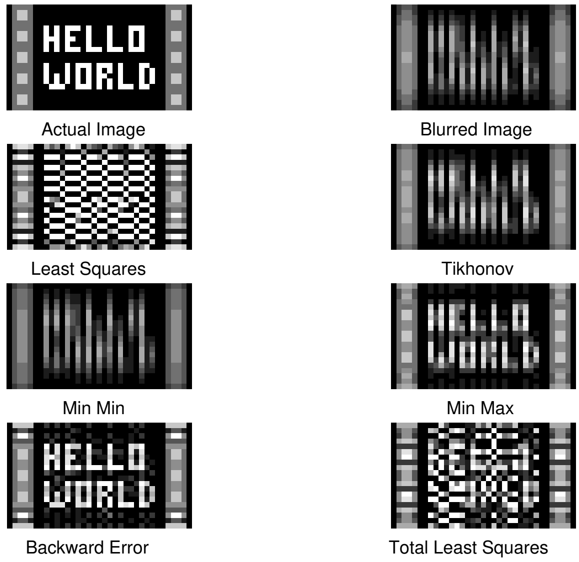

Figure 2: Hello World with $\eta = \|\delta A\|_2$

there will not be enough regulation or deregulation to improve the regression. The error bound was selected to be proportional to the 2-norm of the actual error.

Least squares solution is not readable at all, due to its lack of robustness. Tikhonov makes some small gains, but not enough to be useful. The min min solution makes no noticeable improvements. The min max technique generates a readable result. The BE solution is very readable, though obviously not perfect. Total least squares also makes the image worse, though you can almost make out the some rough letters so it is better than least squares in this case.

7. CONCLUSIONS

Several techniques have been proposed to handle uncertainty in estimation. Each has cases where it provides improvement, thus warranting their consideration. The most promising is the backward error technique, which consistently outperforms other techniques when robustness is needed. For many runs of the numerical examples with different error bounds to simulate different assumptions made by the modeler, the BE technique ended up giving reasonable answers for a longer region. It must be noted that fine tuning of the perturbation error bound is very helpful, though, even for robust systems.

References

[1] P.H. Calamai and A.R. Conn. A Stable Algorithm for Solving the Multifacility Location Problem Involving Euclidean Distances. *SIAM Journal on Scientific and Statistical Computing*, 1:512–526, 1980.

[2] S. Chandrasekaran, G. H. Golub, M. Gu, and A. H. Sayed. Parameter Estimation in the Presence of Bounded Data Uncertainties. *SIMAX*, 19(1):235–252, 1998.

[3] S. Chandrasekaran, G. H. Golub, M. Gu, and A. H. Sayed. An Efficient Algorithm for a Bounded Errors-in-Variables Model. *SIMAX*, 20(4):839–859, 1999.

[4] S. Chandrasekaran, M. Gu, A. H. Sayed, and K. E. Schubert. The Degenerate Bounded Errors-In-Variables Model. *SIMAX*, 23(1):138–166, 2001.

[5] F. Cordellier and J.Ch. Giorot. On the Fermat-Weber Problem with Convex Cost Functionals. *Mathematical Programming*, 14:295–311, 1978.

[6] U. Eckhardt. On an Optimization Problem Related to Minimal Surfaces with Obstacles. In R. Bulirsch, W. Oettli, and J. Stoer, editors, *Optimization and Optimal Control*. Springer-Verlag, 1975.

[7] U. Eckhardt. Weber's Problem and Weiszfeld's Algorithm in General Spaces. *Mathematical Programming*, 18:186–196, 1980.

[8] J.W. Eyster, J.A. White, and W.W. Wierwille. On Solving Multifacility Location Problems Using a Hyperboloid Approximation Procedure. *AIIE Transactions*, 5:1–6, 1973.

[9] R. D. Fierro and J. R. Bunch. Colinearity and Total Least Squares. *SIMAX*, 15:1167–1181, 1994.

[10] R. D. Fierro, G. H. Golub, P. C. Hansen, and D. P. O'Leary. Regularization by Truncated Total Least Squares. *SISC*, 18:1223–1241, 1997.

[11] L. El Ghaoui. and H. Lebret. Robust Solutions to Least-Squares Problems with Uncertain Data. *SIMAX*, 18(4):1035–1064, 1997.

[12] G. H. Golub, P. C. Hansen, and D. P. O'Leary. Tikhonov Regularization and Total Least Squares. *SIMAX*, 30(1):185–194, 1999.

[13] G. H. Golub and C. F. Van Loan. An Analysis of the Total Least Squares Problem. *SIMAX*, 17:883–893, 1980.

[14] G. H. Golub and C. F. Van Loan. *Matrix Computations*. The Johns Hopkins University Press, Baltimore, Md, 1996.

[15] S. Van Huffel and J. Vandewalle. *The Total Least Squares Problem: Computational Aspects and Analysis*. SIAM, Philadelphia, PA, 1991.

[16] H.W. Kuhn. A Note on Fermat's Problem. *Mathematical Programming*, 4:98–107, 1973.

[17] R.F. Love. Locating Facilities in Three-dimensional Space by Convex Programming. *Naval Research Logistics Quarterly*, 16:503–516, 1969.

[18] M. L. Overton. A Quadratically Convergent Method for Minimizing a Sum of Euclidean Norms. *Mathematical Programming*, 27:34–63, 1983.

[19] S. Schechter. Minimization of a Convex Function by Relaxation. In J. Abadie, editor, *Integer and Nonlinear Programming*, pages 177–190. North-Holland, 1970.

[20] H. Voss and U. Eckhardt. Linear Convergence of Generalized Weiszfeld's Method. *Computing*, 25:243–251, 1980.

[21] E. Weiszfeld. Sur le point par lequel la somme des distances de n points donnés est minimum. *Tohoku Mathematics Journal*, 43:355–386, 1937.