

August, 1981

A NEW LAD CURVE FITTING ALGORITHM:
SLIGHTLY OVERDETERMINED EQUATION SYSTEMS IN L_1 .

by

E. SENETA*

W.L. STEIGER**

DCS-TR-104

* Department of Mathematical Statistics
Sydney University
Sydney, N.S.W. 2006
AUSTRALIA

** Department of Computer Science
Rutgers University
New Brunswick, N.J. 08903

TITLE: A New LAD Curve-Fitting Algorithm:
Slightly Overdetermined Equation Systems in L_1 .

RUNNING TITLE: A new LAD Curve-Fitting Algorithm

AUTHORS: Eugene Seneta
Department of Mathematical Statistics
Sydney University
Sydney, N.S.W. 2006
Australia

W.L. Steiger¹
Department of Computer Science
Rutgers University
New Brunswick, N.J., 08903

KEY WORDS & PHRASES LAD, LAD Curve-Fitting

AMS(MDS) CLASSIFICATION

1. Research supported in part by a visiting fellowship in the Department of Statistics, The Australian National University. Computations carried out on the Rutgers Laboratory for Computer Science Research DECSYSTEM 20.

ABSTRACT

We present a new algorithm for the discrete LAD curve-fitting problem for n points in $k \leq n$ dimensional space. When k is about $n/3$ it begins to out-perform the best current methods, and the advantage increases with k . The algorithm should be of interest in approximation theory and in robust regression. Moreover our approach may be useful in linear optimization, especially in linear programming.

I. INTRODUCTION AND SUMMARY

Given n points $(x_i, y_i) \in \mathbb{R}^{k+1}$, the discrete L_1 , or least absolute deviation (LAD) curve-fitting problem seeks $\underline{c} \in \mathbb{R}^k$ to minimize the distance function

$$(1) \quad f(\underline{c}) = \|\underline{y} - X\underline{c}\|_1 = \sum_{i=1}^n \left| y_i - \sum_{j=1}^k x_{ij} c_j \right|$$

where, under obvious notation, we have written $\underline{y}' = (y_1, \dots, y_n)$ and the i^{th} row of X is \underline{x}_i . The minimizing \underline{c} determines a k dimensional hyperplane $\pi = \{(x, y) \in \mathbb{R}^{k+1} : y = \langle \underline{c}, \underline{x} \rangle\}$ that "best" fits the n points in the sense of discrete L_1 .

The LAD fitting method seems to have greater antiquity than least squares. In 1789 Laplace (see Eisenhart [11]) developed an algorithm to minimize (1) when $k=1$. It was stimulated by a geometric method of Boscovich [11] when $\sum_{i=1}^n y_i = c_1 + c_2 \sum_{i=1}^n x_i$; this is the case $k=2$ with $x_{i1} = 1$ and the fit constrained to pass through the centroid of the data.

In 1809 Gauss (see Sheynin [19], p.311) stated without proof that a solution \underline{c} to (1) satisfies $y_i = \langle \underline{c}, \underline{x}_i \rangle$ for a certain set of k points (x_i, y_i) . This fact would imply a crude algorithm that minimizes f by searching all the $\binom{n}{k}$ \underline{c} 's determined by k of the n given points.

Much later Edgeworth [9], [10] formulated a general procedure for (1) and in 1930 Rhodes [15] devised a method for $k=3$ that seems easy to generalize, as pointed out by Singleton [20]. None of these methods was implemented when computers began to appear, possibly because they were too complicated. In fact it may be the computational complexity of LAD (as well as the analytical intractability in a statistical setting) that forced it to take the back seat to least squares in linear regression.

The connection with linear programming (LP) [7], [12], [23] and the development of Simplex codes led to the first automatic procedures for LAD curve fits. Subsequently there have been many contributions to the subject, among them [1], [8], [13], [14], [17], and [22]. In addition the 1977 Communications in Statistics, Vol. B6, # 4, is devoted to LAD and contains a comprehensive reference list.

Three good LAD algorithms have been implemented. Barrodale and Roberts (BR) [3], [4] brought about an important advance by showing how to implement an LP version of (1) in an efficient way by modifying the Simplex method. Later, Bartels, Conn and Sinclair (BCS) [5] used a gradient method to derive what turns out to be a very similar algorithm. Finally Bloomfield and Steiger (BS) [6] have proposed yet another method that reduces the complexity of minimizing (1) even further. All three fit into the general framework laid out by Roberts and ben Israel [16], BR even matching up in the detailed computations.

Anderson and Steiger [2] discuss the similarities and differences between the three methods and study the computational complexities. It is shown that BCS and BS are linear in n and that all three are linear in k ; as n increases, the complexity of BR increases faster than linearly.

In the present paper we address the problem of limiting the computational cost as k increases towards n . We convert (1) into an equivalent problem, one that becomes easier to solve as k grows. Specifically, given y and X , the problem of minimizing f in (1) may be written as

$$(2) \quad \min \|r\|_1 : r = y - Xc, \quad c \in R^k.$$

Suppose X has p independent columns, $p \leq k$. If A is an $n - p$ by n matrix whose rows are orthogonal to the span of the columns of X (i.e., AX an $n - p \times k$ zero matrix), $A\underline{r} = A\underline{y}$, and we write $A\underline{y} = \underline{b} \in \mathbb{R}^{n-p}$. If the rows of A are also linearly independent, (2) is actually equivalent to

$$(3) \quad \min g(\underline{r}) : A\underline{r} = \underline{b}, \text{ where } g(\underline{r}) = \|\underline{r}\|_1$$

This follows from the fact that for any $\underline{c} \in \mathbb{R}^k$, if $\underline{r} = \underline{y} - X\underline{c}$, $A\underline{r} = A\underline{y}$, while if $\underline{r} \in \mathbb{R}^n$ satisfies $A(\underline{r} - \underline{y}) = 0$ $\underline{r} - \underline{y}$ is in the (p dimensional) orthogonal complement of the row space of A (so $\underline{r} = \underline{y} - X\underline{c}$ for some $\underline{c} \in \mathbb{R}^k$).

The correspondence provides further information that is relevant to iterative algorithms for (2) and (3). As Gauss observed, an \underline{r} that minimizes (2) may be chosen so that at least k of its components are zero. The same property also characterizes solutions to (3).

Iterative algorithms for (1) and (2) exploit the characterization and move to the optimum via a sequence $\underline{r}_1, \underline{r}_2, \dots, \underline{r}_N$ of approximations for which $g(\underline{r}_s) \geq g(\underline{r}_{s+1})$ and at least k components of \underline{r}_s are zero. Accordingly, if the j^{th} component of \underline{r}_s is zero for $j \in B = \{j_1, \dots, j_k\}$, the \underline{c} that minimizes f is the solution of the system $\sum_{j=1}^k c_j x_{ij} = y_i, i \in B$. The non-zero components of the optimal \underline{r} satisfy $r_i = y_i - \sum_{j=1}^k c_j x_{ij}, i \notin B$. However if \underline{c} has not been computed, these r_i 's are also the solutions to

$$\sum_{j \notin B} a_{ij} r_j = b_i$$

The problem (3) is of interest in its own right. Seneta [18] points out that it was introduced in a statistical setting by Cauchy.

It is worth noting that (3) is not dual to (1), at least in the sense used in linear programming. It would rather seem that (3) is a problem that is complementary to (1).

Writing the LAD problem in parametric form (1), requires the $n \times (k+1)$ augmented matrix $(X|\underline{y})$. If instead it is written in non-parametric form, (3), the augmented matrix $(A|\underline{b})$ of size $n-p \times n+1$ is required. Assuming $p=k$, the latter is smaller when $k \geq n/2$.

One might therefore expect a good algorithm for (3) to out-perform the best procedures for (1) as $k \geq n/2$ increases, even taking account of the cost in finding A and \underline{b} . This turns out to be the case but surprisingly, the cross over point seems to be about $k = n/3$.

Our algorithm for (3) is based on BS but other LAD procedures could be adapted in a similar way. The idea of minimizing (1) via (3) seems to carry over to functions other than the discrete L_1 norm. For example algorithms for

$$(4) \quad \text{opt } h(\underline{x}) : \underline{x} = \underline{y} - X\underline{c}, \underline{c} \in R^k, h:R^n \rightarrow R \text{ nonlinear}$$

might be simplified using the foregoing ideas, but we do not address this issue here.

Finally (1) via (3) has bearing on bounded, feasible LP problems. The problem

$$(5) \quad \begin{aligned} &\max \langle \underline{d}, \underline{x} \rangle \\ &\text{subj to } \begin{cases} C\underline{x} = \underline{a} \\ \underline{x} \geq \underline{0} \end{cases} \end{aligned}$$

where C is m by n , $m < n$, is equivalent (see [5] or [21]) to a LAD fit for $n+1$ points in R^{m+1} . In non-parametric form, (3), it would use a data structure of size $n-m+1$ by $n+2$, assuming the row rank of C is m , and is smaller than (5) when $m \geq (n+1)^2/(2n+3)$, or

about $n/2$.

In section 2 we give the details of the algorithm for (3), including the passage from (1) to (3). Section 3 contains the results of computational experiments that compare the performance of the new algorithm to BS.

2. THE NEW ALGORITHM

Suppose $A' = (A|b)$ is given, A an m by $n > m$ matrix of full rank and $b \in R^m$. The algorithm we give for (3) is, except for some details, quite similar to BS.

As an initial solution take $\underline{r} = \begin{pmatrix} \underline{z} \\ \underline{o} \end{pmatrix}$, $\underline{o} \in R^{n-m}$ a zero vector, and write $A = (B|N)$, B the $m \times m$ matrix consisting of the first m columns of A , assumed without loss of generality to be linearly independent, N the remaining $n-m$ columns. B is the initial basis and corresponds to the m (generally) non-zero elements of \underline{r} . Because $A \underline{r} = \underline{b}$
 $\underline{z} = B^{-1} \underline{b}$.

To continue, a column, of B , say the p^{th} , will leave the basis via an exchange with a column of N , say the q^{th} , that will enter: z_p will become 0 and o_q will become $t \neq 0$. Our method will select q in a heuristic fashion and, once it is determined, p will be optimally chosen.

Assuming that n_q has been chosen to enter, the next approximation will be found as a member of the one-parameter family

$$\underline{r}'(t) = \begin{pmatrix} \underline{z}'(t) \\ \underline{o} \end{pmatrix} + t \underline{e}_{-m+q}$$

where $\underline{e}_i \in R^n$ is the i^{th} unit coordinate vector and $\underline{z}'(0) = \underline{z}$. Clearly $\underline{r}'(0)$ is the current solution.

Since $A \underline{r}'(t) = \underline{b}$, $B \underline{z}'(t) + t \underline{n}_q = \underline{b}$, \underline{n}_q denoting the q^{th} column of N . Thus

$$\begin{aligned}
 \underline{z}'(t) &= B^{-1}\underline{b} - t B^{-1}\underline{n}_q \\
 (6) \qquad &= \underline{z} - t \underline{v}
 \end{aligned}$$

where we write $\underline{v} = B^{-1}\underline{n}_q$. The current value of the objective g in (3) will now be

$$\begin{aligned}
 g(\underline{r}'(t)) &= \sum_{i=1}^n |r'_i(t)| \\
 (7) \qquad &= \sum_{i=1}^m |z'_i(t)| + |t| \\
 &= \sum_{i=1}^m |z_i - t v_i| + |t|
 \end{aligned}$$

The value of t that minimizes (7) defines a line $y=tx$ that is the LAD fit through the origin for the points (v_i, z_i) , $i=1, \dots, m$ and $(1,0)$. It is easily recognized to be the weighted median \hat{t} of the ratios z_i/v_i , with weight $|v_i|$, and 0, with weight 1. (see [2] or [6]).

If $\hat{t} = 0$ \underline{n}_q may not enter the basis so another column of N would be investigated. Otherwise $\hat{t} = z_p/v_p$, for some p , $1 \leq p \leq m$, and the p^{th} term in the sum in (7) becomes zero. This means that $z'_p(\hat{t}) = 0$ so that \underline{n}_q replaces \underline{b}_p (the p^{th} column of B) in the next basis. The objective in (7) cannot increase because when $t = 0$, $\underline{r}'(t) = \underline{r}$ is the current solution and $g(\underline{r}'(0)) \geq g(\underline{r}'(\hat{t}))$.

We now describe the choice of \underline{n}_q , the non-basic column that will enter. Ideally one seeks that column which, when it optimally replaces some basis column, will produce the greatest reduction in the objective (3). A brute force "look ahead" method would compute $\underline{v} = B^{-1}\underline{n}_q$ for each column \underline{n}_q of N , minimize (7), and then enter that column corresponding

to the smallest value of $g(r'(\hat{t}))$. Our heuristic method, based on that of BS, avoids the computational cost of actually minimizing (7) for each \underline{n}_q , and usually chooses the best column anyway.

The heuristic is based on weighted medians. If \underline{n}_q were chosen, the foregoing procedure for optimally selecting a column to leave the basis would compute $\underline{v} = B^{-1} \underline{n}_q$ and minimize (7). The minimum occurs at \hat{t} , the weighted median of the ratio z_i/v_i with weights $|v_i|$, and 0 with weight 1. We assume $\hat{t} \neq 0$ or else \underline{n}_q cannot enter. One can also express \hat{t} as the median of the distribution function

$$(8) \quad F(t) = (I_0(t) + \sum_A |v_i|) / (1 + \sum_{i=1}^m |v_i|),$$

where $I_0(t)$ is 1 on $[0, \infty)$ and 0 otherwise, and $A = \{i : z_i/v_i \leq t\}$. By definition, $F(\hat{t}) \geq 1/2$ and $F(t) < 1/2$ for $t < \hat{t}$.

In (6) the current approximation corresponds to $t=0$, and the next approximation to \hat{t} . For this reason the quantity

$$(9) \quad |1/2 - F(0)|,$$

approximately $|F(\hat{t}) - F(0)|$, is a rough measure of the distance - in terms of F - between the current approximation and the next one. It purports to capture the relative advantage of using \underline{n}_q to optimally replace a basic column.

If in (8) $F(0) = 1/2$, replacement with \underline{n}_q would not decrease g in (7). Excluding this case, either $F(0) < 1/2$ or $F(t) \rightarrow u > 1/2$ as $t \uparrow 0$, by the right continuity of F . The criterion in (9) would then become

$$(10) \quad \begin{array}{ll} 1/2 - F(0) & \text{if } F(0) < 1/2 \\ F(0^-) - 1/2 & \text{if } F(0^-) > 1/2 \end{array}$$

From (8), $1/2 - F(0) = 1/2 - [(1 + \sum_n |v_i| + \sum_z |v_i|)] / (1 + \sum_{i=1}^m |v_i|)$

which simplifies to

$$\frac{-1 - \frac{\sum_n |v_i|}{n} - \frac{\sum_z |v_i|}{z} + \frac{\sum_p |v_i|}{p}}{2(1 + \sum_{i=1}^m |v_i|)}$$

where n denotes the indices for which $z_i/v_i < 0$, z the indices for which $z_i=0$, and p , those for which $z_i/v_i > 0$.

Similarly $F(0^-) - 1/2$ may be written as

$$(12) \quad \frac{-1 + \frac{\sum_n |v_i|}{n} - \frac{\sum_z |v_i|}{z} - \frac{\sum_p |v_i|}{p}}{2(1 + \sum_{i=1}^m |v_i|)}$$

so the criterion in (10) may be expressed as

$$(13) \quad c_q = \frac{\left| \frac{\sum_n |v_i|}{n} - \frac{\sum_z |v_i|}{z} \right| - 1 - \frac{\sum_p |v_i|}{p}}{2(1 + \sum_{i=1}^m |v_i|)}$$

It would be evaluated for each column \underline{n}_q of N . The one with the largest value of $c_q > 0$ is chosen to enter.

If $c_q \leq 0$ for all columns in N the algorithm terminates at an optimal \underline{r} , unless the current approximation is degenerate and $r_i=0$ for $j > n-m$ residuals (so one of the $z_i=0$). This may be seen by interpreting the numerator of (13) as a directional derivative in the direction from \underline{r}' to \underline{r} . Minus the right hand derivative of (7) at $t=0$ is the numerator of the expression in (11) which therefore measures the rate of decrease in g as one moves away from the current approximation in the direction $\underline{w}_q = \underline{r}'(1) - \underline{r}'(0) = \begin{pmatrix} z \\ 0 \end{pmatrix} - \begin{pmatrix} z-v \\ 0 \end{pmatrix} + \underline{e}_{m+q} - \begin{pmatrix} z \\ 0 \end{pmatrix} = \begin{pmatrix} -v \\ 0 \end{pmatrix} + \underline{e}_{m+q}$, where $\underline{v} = B^{-1} \underline{n}_q$. Similarly the left hand derivative at $t=0$ in (7) is the numerator of (12), so it measures the rate of decrease as one moves

from the current approximation in the direction $-w_q$.

If $c_q \leq 0$, for a certain q , the directional derivation of g along the line $\underline{r} + t w_q \in \Pi = \{x : Ax = b\}$ is always non-negative, by convexity, and $g(\underline{r}) \leq g(\underline{r} + t w_q)$, all t . If $c_q \leq 0$ for all $q=1, \dots, n-m$, $g(\underline{r}) \leq g(\underline{r} + t w_q)$ for all the w_q so no non-basic column may replace a basic one and reduce g .

If no basic residual is zero, the current \underline{r} is optimal, by the geometry of the problem. The graph of g is a convex polytope in R^{n+1} comprised of $n-m$ dimensional hyperplanar "pieces". The vertices correspond to points $(\underline{r}, g(\underline{r}))$ for which $A\underline{r} = \underline{b}$ (an $n-m$ dimensional hyperplane in R^n) intersected with $n-m$ hyperplanes $r_{i_1} = 0, r_{i_2} = 0, \dots, r_{i_{n-m}} = 0$, for any choice of $n-m$ different indices from the set $\{1, \dots, n\}$. Such $\underline{r} \in R^n$ are 0-dimensional hyperplanes, and there are $\binom{n}{m}$ of them.

Certain pairs of vertices are connected by edges, 1-dimensional hyperplanes in the polytope defined by $\{(\underline{r}, g(\underline{r})) : A\underline{r} = \underline{b} \text{ and } r_{i_1} = 0, \dots, r_{i_{n-m-1}} = 0\}$. A given vertex would have at least $n-m$ edges emanating from it, one corresponding to the relaxation of each of the conditions $r_{i_1} = 0, \dots, r_{i_{n-m}} = 0$.

A degenerate vertex $(\underline{r}, g(\underline{r}))$ has $j > n-m$ edges, one for each condition $r_i = 0, i=i_1, \dots, i_{n-m}$; $j-n+m$ other r_i 's = 0 whilst maintaining $A\underline{r} = \underline{b}$. In the non-degenerate case the condition $c_p \leq 0$, all p , guarantees that there is no improvement along any edge, and by convexity in any direction in Π .

While it is easy to diagnose degeneracy, it may be complicated to deal with. For this reason we finesse the issue and leave degeneracy as a separate problem, to be treated on its own.

The criterion, (13), for selecting an entering column is analogous to that from the BS LAD algorithm. Other LAD methods like BR and BCS choose the column based only on directional derivatives. BS seems

to be the only method using normalized directional derivatives. This has the nice property of removing scaling differences between the n_q , and may account for the superior performance.

A convenient data structure for implementing this algorithm begins with the augmented matrix $A' = (B|N|b)$ premultiplied by B^{-1} . The matrix

$$(14) \quad D = (I|B^{-1}N|B^{-1}b)$$

contains the information that the vector $\sigma = (\sigma(1), \dots, \sigma(m))$ of indices pointing to the basic columns is $(1, \dots, m)$, that the vector $\sigma^c = (\sigma^c(1), \dots, \sigma^c(n-m))$ of indices pointing to non-basic columns is $(m+1, \dots, n)$, and that $\underline{z} = B^{-1}b$ gives the values of the non-zero components of \underline{r} : $r_{\sigma(i)} = z_i$ and $r_{\sigma^c(i)} = 0$.

To find the column of N that will become basic the criterion c_q in (13) is evaluated for each $q=1, \dots, n-m$. The v 's are the columns of $B^{-1}N$, from (6).

Once the q^{th} element of σ^c (call it j ; initially $j=\sigma^c(q)=m+q$) has been selected to enter, and the p th element of σ is chosen to leave, (14) is updated by a pivot step using Jordan elimination: the p^{th} row of D is multiplied by a constant so the j^{th} column entry is 1; then multiples of row p are added to each other row so that j^{th} column entries become zero.

If E is the m by m matrix that effects these steps, D becomes

$$(15) \quad D' = ED = (E|EB^{-1}N|EB^{-1}b)$$

The new basis pointers are $\sigma' = (1=1, \dots, p-1, m+q, p+1, \dots, m)$ and the non-basic ones are $(\sigma')^c = (m+1, \dots, m+q-1, p, m+q+1, \dots, n)$. EB^{-1} is the inverse of the new basis matrix $B' = A(\sigma')$, formed by taking columns from A according to σ' , and $\underline{z}' = EB^{-1}b$ are the non-zero components of the new solution \underline{r}' . In general, after several steps σ points to m columns of the matrix

that comprise an identity, σ^c to the $n-m$ columns corresponding to $B^{-1}N$ for the current basis/nonbasis partition of A , and the $n+1^{\text{st}}$ column gives the values of the basic residuals : $r_{\sigma(1)} = z_i$.

Suppose we are given X and y in (1), X an n by k matrix of full rank. To convert the LAD fitting problem into (3) we require an

$n-k$ by n matrix of A full row rank that satisfies $AX = 0$. Write $A = (B|N)$, B , $n-k$ by $n-k$, and $X = \begin{bmatrix} X_B \\ X_N \end{bmatrix}$, where X_B is $n-k$ by k and X_N is k by k , and we suppose without loss of generality that X_N is invertible. Then $BX_B + NX_N = 0$ and if we take $B = I$, $NX_N = -X_B$. Transposing, N is defined by

$$(16) \quad X_N^T N^T = -X_B^T,$$

$n-k$ linear systems each of size k . Using Gaussian elimination, and backsolving, N can be found in about $k^3/3 + (n-k)k^2 = nk^2 - 2k^3/3$

multiplication and division operations. This initializes (14) with $(I|N|b)$, where $b = Ay$. It is interesting that an initial solution using BS directly on (1) would cost nk^2 steps, so translation from (1) to (3) seems to require no computational overhead.

This description appears in succinct algorithmic form in the appendix. The performance of the algorithm is compared with that of BS on a variety of LAD problems, (1). The results appear in the next section.

3. COMPUTATIONAL EXPERIENCE

In this section we compare the computational cost of solving (1) directly using BS with that of translating (1) to (3), solving (3) using the algorithm (SS) of the previous section, and then obtaining the optimal c . Based on some simple assumptions, SS should be cheaper, even when $k < n/2$, a conjecture that is supported by actual experience.

One iteration of BS costs (in terms of multiplications and divisions) $n(k+1)$ operations to update, plus the cost, $WM(n-k)$ of forming $n-k$ ratios and then finding the weighted median. The first k iterations successively pivot new points into the fit and may be regarded as a start up phase, at total cost of $k[n(k+1)+WM(n-k)]$. If N further iterations are required the total work for BS would be

$$(17) \quad k[(n(k+1) + WM(n-k))] + N[n(k+1) + WM(n-k)] + \frac{k^3}{3} + k^2,$$

the last two terms reflecting the extra cost of solving $\underline{y} = X\underline{c}$ for \underline{c} , once it is known which k rows of X determine the optimal fit.

To initialize SS for (3) it costs $nk^2 - 2k^3/3$ steps to find A plus a further $(n-k)k$ to compute $\underline{b} = A\underline{y}$. Each iteration then costs $(n+1)(n-k)+WM(n-k)$, the first term being the cost of updating D in (14). Thus if M iterations of SS are required, the total work would be about

$$(18) \quad nk^2 + nk - 2k^3/3 - k^2 + M[(n+1)(n-k)+WM(n-k)] + \frac{(n-k)^3}{3} + (n-k)^2$$

steps, the last two terms being the extra cost of finding \underline{c} once the optimal \underline{r} is known.

The difference between the BS and SS for just the start up phase and the final solve is therefore

$$(19) \quad k WM(n-k) + \frac{2k^3}{3} + k^2 + \left[\frac{k^3 - (n-k)^3}{3} \right] + [k^2 - (n-k)^2].$$

Since $WM(n-k)$ is $n-k$ divisions for the ratios, plus the weighted median calculation, (19) is positive even if k is substantially less than $n-k$.

On the other hand if N and M (the non-startup iterations required by BS and SS respectively) are comparable, that phase of SS will be cheaper than the corresponding one for BS when $k \geq n/2$. Putting these observations together, one expects SS to outperform BS even when $k < n/2$, and for fixed n , the advantage should increase with k .

To evaluate the validity of the assumptions on which the foregoing argument is based, we performed actual comparisons of BS with SS. First we approximated $f(x) = \sqrt{x}$ on $[0,1]$ by the LAD polynomial ϕ_{k-1} of degree $k-1$, based on $n+1$ equally spaced points $t_i = i/n$, $i = 0,1,\dots,n$. We thus minimized

$$\sum_{i=0}^n \left| \sqrt{t_i} - \sum_{j=1}^k c_j t_i^{j-1} \right|,$$

a fit of the form (1) with $y_i = \sqrt{t_i}$ and $x_{ij} = t_i^{j-1}$. We used $n=15$ and $k=5,7,9,11$.

The results appear in Table 1. The iteration counts for BS include the k startup steps, so they measure $N+k$. These results do not seem to contradict the assumption that M from (18) and N (from (17)) are comparable. Execution timings may be misleading because they reflect features of hardware, and peculiarities of coding. Nevertheless the figures in Table 1 do accord with the conjecture that SS is more efficient than BS when $k > n/3$ and the advantage improves as k increases.

Table 1. Comparison of CPU times and iteration counts for LAD approximation of \sqrt{x} on $[0,1]$ by $\sum_{j=1}^k c_j t^{j-1}$ based on 16 equally spaced points $t_i = i/15$, $i=0,\dots,15$.

	k=5		k=7		k=9		k=11	
	CPU	ITER	CPU	ITER	CPU	ITER	CPU	ITER
BS	.036	11	.044	12	.082	16	.088	12
SS	.032	7	.042	7	.046	9	.042	6

Next we fit regression models of the form

$$(20) \quad Y = c_1 X_1 + \dots + c_k X_k + u$$

following the procedure of [2] . Thus a sample of size n from (20) was generated as follows. For each $i=1, \dots, n$, successive random numbers $x_{i2}, x_{i3}, \dots, x_{ik}$, u_i were generated. We set $x_{i1} = 1$ and

formed $y_i = c_1 x_{i1} + \dots + c_k x_{ik} + u_i$. The c_j were taken to be \sqrt{j} so the columns of $X = (x_{ij})$ would be scaled differently.

The x_{ij} and u_i were taken to be centered Pareto random numbers of index $\alpha = 1.2$, generated using the density function

$$P(t) = \alpha / [1 + (t-a)^{1+\alpha}] , \quad t \geq a = -\alpha / (\alpha - 1) .$$

They would have mean zero and infinite variance.

Once $(X|y)$ is generated LAD estimates of c using BS and SS are obtained. The process is then repeated independently until 10 samples from (20) have been considered. In this way the Monte-Carlo results won't be too strongly influenced by a short sequence from the random number generator.

Table 2 contains results for $n=10$ and Table 3 for $n=50$. The CPU times are net of generating the samples of (20) and are accumulated over the 10 repetitions; the iteration counts are also summed.

In both tables, the iterations differ by about 10k. Since BS iterations include the k startup steps, this finding again supports the assumption that M from (18) and N from (17) are comparable, since over 10 repetitions we would expect such a difference between the counts. Also the CPU timings suggest that $SS < BS$ for $k > n/3$ and that the difference, $BS - SS$ increases with k .

Table 2 Comparison of CPU times and iteration counts summed over 10 independent repetitions of fitting the model (20) for $n=10$ and X, U distributed as Pareto (1.2)

	k=4		k=6		k=8	
	CPU	ITER	CPU	ITER	CPU	ITER
BS	.16	65	.28	86	.32	94
SS	.18	31	.16	25	.12	10

Table 3 Comparison of CPU times and iteration counts summed over 10 independent repetitions of fitting the model (20) for $n=50$ and X, U distributed as Pareto (1.2)

	k=18		k=22		k=26		k=30		k=34	
	CPU	ITER	CPU	ITER	CPU	ITER	CPU	ITER	CPU	ITER
BS	7.46	324	10.04	362	13.06	395	17.20	452	20.70	473
SS	6.50	180	7.16	189	7.34	166	7.92	170	8.46	158

Acknowledgement:

We wish to thank Peter Bloomfield for valuable discussions on the equivalence of (1) and (3).

REFERENCES

- [1] Abdelmalek, N.N. (1971). "Linear L_1 Approximation for a Discrete Point Set and L_1 Solutions of Overdetermined Equations", JACM 18, 41-47
- [2] Anderson, D. and Steiger, W.L. (1980). "A comparison of Methods for Discrete L_1 Curve Fitting", Department of Computer Science Technical Report, #96, Rutgers University.
- [3] Barrodale, I. and Roberts, F.D.K. (1973). "An Improved Algorithm for Discrete L_1 Linear Approximation", SIAM J. Numer. Anal. 10, 839-848
- [4] Barrodale, I. and Roberts, F.D.K. (1974). "Algorithm 478: Solution of an Overdetermined System of Equations in the L_1 Norm", Comm. A.C.M. 14, 319-320
- [5] Bartels, R., Conn, A. and Sinclair, James W. (1978). "Minimization Techniques for Piecewise Differentiable Functions: the L_1 Solution to an Overdetermined System", SIAM J. Numer. Anal. 15, 224-241
- [6] Bloomfield, P. and Steiger, W.L. (1980). "Least Absolute Deviations Curve-Fitting", SIAM J. Scientific and Statistical Comp. 1, 290-301.
- [7] Charnes, A., Cooper, W.W. and Ferguson, R.O. (1955). "Optimal Estimation of Executive Compensation by Linear Programming", Mgt. Sci. 1, 138-151
- [8] Claerbaut, Jon and Muir, Francis (1973). "Robust Modeling with Erratic Data", Geophysics 38, 826-844
- [9] Edgeworth, F.Y. (1887). "A New Method of Reducing Observations Relating to Several Quantities", Phil. Mag. (Fifth Series) 24, 222-223
- [10] Edgeworth, F.Y. (1888). "On a New Method of Reducing Observations Relating to Several Quantities", Phil. Mag. (Fifth Series) 25, 184-191
- [11] Eisenhart, C. (1961). "Boscovich and the Combination of Observations" In Roger Joseph Boscovich, L.L. Whyte, Ed. New York: Fordham University Press
- [12] Harris, T. (1950). "Regression Using Minimum Absolute Deviations", Amer. Statistician 4, 14-15
- [13] Karst, Otto J. (1958). "Linear Curve Fitting Using Least Deviations", J. Amer. Stat. Assoc 53, 118-132
- [14] McCormick, G.F. and Sposito, V.A. (1976). "Using the L_2 Estimator in L_1 Estimation", SIAM J. Numer. Anal. 13, 337-343

- [15] Rhodes, E.C. (1930). "Reducing Observations by the Method of Minimum Deviations", Phil. Mag. (Seventh Series) 9, 974-992
- [16] Robers, P. and ben Israel, A. (1970) "A Suboptimization Method for Interval Linear Programming : A New Method for Linear Programming", Linear Algebra and its Applications 3, 383-405.
- [17] Schlossmacher, E.J. (1973). "An Iterative Technique for Absolute Deviations Curve Fitting", J.Amer.Stat.Assoc. 68, 857-865.
- [18] Seneta, E. (1976) "On a Contribution of Cauchy to Linear Regression Theory", Ann. de la Société Scientifique de Bruxelles 90, 229-235.
- [19] Sheynin, O.B. (1973). "R.J. Boscovich's Work on Probability" Archive for History of Exact Sciences, 9, 306-324.
- [20] Singleton, R.R. (1940). "A Method for Minimizing the Sum of Absolute Values of Deviations", Ann.Math.Stat. 11, 301-310.
- [21] Steiger, W.L. (1980). "Linear Programming via Discrete L_1 Curve Fitting". Dept. of Computer Science Technical Report, #97, Rutgers University.
- [22] Usow, K.H. (1967). "On L_1 Approximation II : Computation for Discrete Functions and Discretization Effects", SIAM J. Numer. Anal. 4, 233-244.
- [23] Wagner, Harvey, M. (1959). "Linear Programming Techniques for Regression Analysis", J. Amer. Stat. Assoc. 54, 206-212.

APPENDIX

Given $(x_i, y_i) \in \mathbb{R}^{k+1}$, $i=1, \dots, n$ as in (1), write
 $m = n - k$ and $X = \begin{pmatrix} x_1 \\ \vdots \\ x_n \end{pmatrix}$. The algorithm SS may be described as follows.
Initialize:

(1) Renumbering the rows of $(X|y)$ if necessary so that X_N , the bottom k rows of X , is invertible, solve the m linear systems $X_N^T N = -X_B^T$ for N , where X_B denotes the top m rows of X .

(2) $D \leftarrow (A|b)$ where $A \leftarrow (I|N)$ is m by n and $b = Ay$.

(3) $\sigma \leftarrow (1, \dots, m)$, $\sigma^c \leftarrow (m+1, \dots, n)$.

(4) $r_{\sigma(i)} \leftarrow b_i$, $i=1, \dots, m$, $r_{\sigma^c(i)} \leftarrow 0$, $i=m+1, \dots, n$.

Choose Entering Column

(5) $z \leftarrow \{i, 1 \leq i \leq m : b_i = 0\}$

(6) For $j=1$ thru k DO

$$v_i \leftarrow D_{i \sigma^c(j)}, \quad i=1, \dots, m$$

$$n \leftarrow \{i : b_i, v_i \text{ opposite sign}\}$$

$$p \leftarrow \{1, \dots, m\} \setminus n \setminus z$$

$$c_j \leftarrow \frac{\left| \frac{\sum_n |v_i|}{n} - \frac{\sum_p |v_i|}{p} \right| - 1 - \frac{\sum_z |v_i|}{z}}{1 + \sum_i |v_i|}$$

END

(7) $S \leftarrow \{1, \dots, k\}$

(8) $c_q \leftarrow \max_S c_j$

- (9) If $c_q > 0$ go to 11
- (10) If $\prod_{i=1}^m r_{\sigma(i)} = 0$ Return "degenerate",
 else solve $y_{\sigma^c(i)} = \sum_{j=1}^k x_{\sigma^c(i)j} \theta_j$ for θ
 and STOP.

Find Leaving Column.

- (11) $v_i \leftarrow D_{i, \sigma^c(q)}$, $i=1, \dots, m$
- (12) $\hat{t} \leftarrow$ weighted median of
 $b_1/v_1, \dots, b_m/v_m, 0$ with weights $|v_1|, \dots, |v_m|, 1$.
- (13) If $\hat{t} = b_p/v_p \neq 0$, go to 16
- (14) $S \leftarrow S \setminus \{q\}$; If $S = \emptyset$, go to 10.
- (15) Go to 8.

Update

- (16) divide row p of D by $D_{p \sigma^c(q)}$
- (17) For $i \neq p$, (row i of D) \leftarrow (row i of D)
 $-$ (row p of D) $\cdot D_{i \sigma^c(q)}$
- (18) $\sigma(p) \leftrightarrow \sigma^c(q)$
- (19) $r_{\sigma(i)} \leftarrow b_i$ $i=1, \dots, m$; $r_{\sigma^c(q)} \leftarrow 0$
- (20) Go to 5