ON THE PROBLEM OF OBTAINING NUMERICAL SOLUTIONS TO LEAST SQUARES EQUATIONS

BU-222-M E. C. TOWNSEND June 1966

ABSTRACT

This paper discusses some of the difficulties which may be encountered when obtaining numerical solutions to least squares equations. The difficulties are illustrated by means of a pathological example and a hypothetical computer.

Biometrics Unit, Plant Breeding Department, Cornell University, Ithaca, N. Y.

ON THE PROBLEM OF OBTAINING NUMERICAL SOLUTIONS TO LEAST SQUARES EQUATIONS

BU-222-M E.C. TOWNSEND June 1966

Introduction

A task frequently encountered by experimenters is that of solving the least squares normal equations associated with the linear model

$$y = Xb+e.$$

In this model y_{nx1} is a vector of observations, X_{nxk} is a known matrix of coefficients, b_{kx1} is a vector of parameters to be estimated and e_{nx1} is a vector of random residuals having zero expectation. The "normal equations" $X'X\hat{b} = X'y$ when $X'X$ is of rank k are solved as $\hat{b} = (X'X)^{-1}X'y$. The computation of $X'X$, of its inverse, of $X'y$ and then of $\hat{b}$ is a laborious task for large n that is quickly relegated to a computer if one is available. A program (set of instructions for the computer) is developed to perform these tasks and may be used repeatedly for long periods of time with excellent results. However, it is possible that a lengthy period of satisfactory usage may develop in the user a false sense of confidence in the machine computations.

A recent cartoon showed the motto "To err is unlikely, to forgive unnecessary" displayed in a computing center. Unfortunately, in the minds of many users this is no joke, because the faith displayed by some people in computer output could be characterized by "To err is impossible." Such users may be bewildered when, after submitting what appears to be a very ordinary least squares problem to the computer, they obtain unreasonable or even impossible solutions. This situation may arise either because of the nature of the equations, or because of the structure of the machine or, more commonly, a combination of both. These matters are now discussed.

Biometrics Unit, Plant Breeding Department, Cornell University, Ithaca, N. Y.

Example

Suppose we have

$$y = \begin{bmatrix} 0.5 \\ 0.1 \\ 1.0 \\ 0.0 \\ 0.1 \end{bmatrix} \qquad \text{and } X = \begin{bmatrix} 0.1 & 0.1 \\ 0.9 & 1.1 \\ 0.4 & 0.5 \\ 0.5 & 0.6 \\ 0.9 & 1.1 \end{bmatrix} .$$

Then, by direct multiplication

$$X'X = \begin{bmatrix} 2.04 & 2.49 \\ 2.49 & 3.04 \end{bmatrix} \qquad \text{and } X'y = \begin{bmatrix} 0.45 \\ 0.55 \end{bmatrix} .$$

Hence the normal equations are

$$\begin{bmatrix} 2.04 & 2.49 \\ 2.49 & 3.04 \end{bmatrix} \hat{b} = \begin{bmatrix} 0.45 \\ 0.55 \end{bmatrix} . \qquad (1)$$

It is readily seen that $\hat{b}' = (-1.0, 1.0)$ satisfies these equations; i.e. the exact solution to this set of equations is

$$\hat{b} = (X'X)^{-1}X'y = \begin{bmatrix} -1.0 \\ 1.0 \end{bmatrix} . \qquad (2)$$

The nature of the equations

Linear equations may be "ill-conditioned", a term used to describe equations in which very small changes in their structure, or in the solving process, result in large changes in the solution. A common measure of ill-conditioning of linear equations is the determinant of the coefficient matrix--the smaller the determinant the more ill-conditioned the equations.

Using the example given above, if $X'y$ had been $\begin{bmatrix} .451 \\ .550 \end{bmatrix}$ instead of $\begin{bmatrix} .45 \\ .55 \end{bmatrix}$ then $\hat{b}$ would have been, correct to 2 decimal places, $\begin{bmatrix} 1.03+ \\ -.66 \end{bmatrix}$ instead of the $\begin{bmatrix} -1.0 \\ 1.0 \end{bmatrix}$ obtained in (2). In this case a very small change in one element of the equations has had a very great effect on the solution. In general then, we should always look at the design matrix of an experiment to see that it does not yield an ill-conditioned coefficient matrix; for if it does, the solution $\hat{b}$ to the normal equations can be influenced greatly by small errors in the observations $\hat{y}$; and this should be avoided whenever possible.

Notice, the problem illustrated here of large changes in a solution arising from small changes in the right-hand side is totally unrelated to using computers to obtain solutions. We now come to difficulties encountered as a result of using digital computers to solve least squares equations.

Memory structure of a computer

Most computer storage units ("memory") are divided into subdivisions called "words". Each word is of finite size capable of representing varying sizes of numbers, depending on the model of computer. Usually one word is used to store one number, although with special programming techniques more than one word may be used to store a number. For purposes of discussion, assume the machine being used has word size of 48 binary bits or positions. In "fixed point arithmetic" one of these positions is used to record the sign, while the remaining 47 contain a binary number; in total these correspond to approximately 14 decimal digits.

Using storage of the nature just described requires keeping track of decimal points in programming. The most common procedure is to use "floating point arithmetic" in which a number is expressed as a fractional part of the appropriate power of 10. For example, 24 may be expressed as $.24 \times 10^3$ and $.0024$ as $.24 \times 10^{-2}$.

To accommodate this information a 48 binary position word may be segmented as follows:

Contents	No. of Bits
Sign of the fractional part	1
Sign of the exponent 10	1
Exponent	10
Fractional part	36

Using this arrangement, a maximum of 11 significant digits of a number may be stored.

Difficulties that can arise from finite word size

Finite word size may create two problems: 1) the true members needed in a calculation may not be presentable to the machine, and 2) rounding error during computation may become great enough to give erroneous answers. To illustrate these problems we will imagine using a hypothetical computer capable of storing only two significant digits in floating point arithmetic. Consider solving equations (1) with this machine, namely

$$\begin{bmatrix} 2.04 & 2.49 \\ 2.49 & 3.04 \end{bmatrix} \hat{b} = \begin{bmatrix} 0.45 \\ 0.55 \end{bmatrix}$$

As they stand, these equations cannot be given to our computer because the elements of X'X have three significant digits but the machine can only store two. Therefore an approximation to the equations has to be submitted for solution; suppose the approximation is

$$\begin{bmatrix} 2.0 & 2.5 \\ 2.5 & 3.0 \end{bmatrix} \hat{b} = \begin{bmatrix} 0.45 \\ 0.55 \end{bmatrix} . \quad (3)$$

The true solution to this set of equations is

$$\hat{b} = \begin{bmatrix} 0.1 \\ 0.1 \end{bmatrix} . \quad (4)$$

Observe that the solution to these equations, which approximate those of real interest, namely (1), has no apparent relationship to the solution of those original equations. This change in $\hat{b}$ results solely from the numbers involved in (1) being limited to two significant digits, in accord with the limited word size of our hypothetical computer. This illustrates how a computer, though its finite word size, may so seriously affect our ability to give it an exact problem that the solution to an approximation to that problem may in no way resemble the solution to the true problem.

The inability to record the results of computations precisely due to finite word size is called rounding error. Continuing the example, in the approximate equations, (3), the true inverse of

$$\begin{bmatrix} 2.0 & 2.5 \\ 2.5 & 3.0 \end{bmatrix} \text{ is } \frac{1}{2.0 \times 3.0 - (2.5)^2} \begin{bmatrix} 3.0 & -2.5 \\ -2.5 & 2.0 \end{bmatrix} = \begin{bmatrix} -12 & 10 \\ 10 & -8 \end{bmatrix} .$$

However, were the computations to be performed in the computer whose words store only two digits, the individual calculations involved in this inverse would be

$2.0 \times 3.0 - (2.5)^2$ computed as $.60 \times 10^1 - .63 \times 10^1 = -.30 \times 10^0 = -.3$

$3.0/-.3$ computed as $-.10 \times 10^2 = -10$

$2.5/-.3$ computed as $.83 \times 10^1 = 8.3$

$2.0/-.3$ computed as $-.67 \times 10^1 = 6.7$

Hence the inverse obtained by the computer would be

$$\begin{bmatrix} -10 & 8.3 \\ 8.3 & 6.7 \end{bmatrix}$$

and so $\hat{b}$ would be obtained as

$\hat{b}_1 = -10(.45) + 8.3(.55)$, in turn computed as $-4.5 + 4.7 = 0.2$ (5)

$\hat{b}_2 = 8.3\ (.45) - 6.7\ (.55)$, in turn computed as $3.7 - 3.7 = 0.0$

Thus we see how rounding error has changed the solution of equations (3) from $\hat{b} = \begin{bmatrix} 0.1 \\ 0.1 \end{bmatrix}$ to $\hat{b} = \begin{bmatrix} 0.2 \\ 0.0 \end{bmatrix}$.

Notice that in this example two kinds of approximation have arisen in trying to solve equations (1). Because of the computer's limited word size (2 significant digits) they were presented to the computer in the form of (3); and again because of word size, equations (3) are solved as (5), although their true solution is (4). Both of these kinds of approximations can arise in real-life problems.

We have just seen how rounding error can have a marked effect on the solution to a set of equations. Rounding error can also make least squares equations impossible to solve. For example if

$$X = \begin{bmatrix} 1 & 1 & 1 & 1 \\ c & 0 & 0 & 0 \\ 0 & c & 0 & 0 \\ 0 & 0 & c & 0 \\ 0 & 0 & 0 & c \end{bmatrix}$$

then

$$X'X = \begin{bmatrix} 1+c^2 & 1 & 1 & 1 \\ 1 & 1+c^2 & 1 & 1 \\ 1 & 1 & 1+c^2 & 1 \\ 1 & 1 & 1 & 1+c^2 \end{bmatrix}.$$

The rank of $X'X$ is 4 and its determinant is $(4 + c^2)c^6$. Suppose now, that $X'X$ was being used in normal equations: the inverse of $X'X$ is required Consider what happens if c is small compared to 1, and $X'X$ is calculated on a digital computer. With c being small, $1 + c^2$ may well be rounded to 1 and then, as far as the computer is concerned, $X'X$ would be of the form

$$X'X = \begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 \end{bmatrix}$$

which is singular. Therefore, for computing purposes, $(X'X)^{-1}$ would not exist.

Frequency of the Problem

The difficulties discussed and illustrated above are not frequently encountered. For the average reader several years of experience may be involved before he is confronted with such difficulties and therefore he should not become unduly alarmed. However, he should be aware of the possibility of these difficulties and he should know under what conditions they are likely to be found.

Caution is advised when analyzing data from experiments which yield a design matrix X whose columns are highly correlated. Under these conditions the determinant of the X'X matrix will be small. Thus the normal equations may be ill-conditioned which can cause the solution to be adversely affected by small errors in the observations y, and by rounding error. Most computer programs used to fit linear models by least squares procedures will compute the correlations between the columns of X if the user requests them. It is recommended that this option be taken and if correlations of .97 or larger are found the user should be aware of the possible ill-conditioning.

When the individual elements of a column of an X matrix have several significant digits and many observations are being used precise representation of the coefficients of the normal equations may not be possible. For example, when a column of X contains elements having five significant digits the squares of the elements will have 10 significant digits. If more than a hundred observations are used the element of X'X corresponding to the sum of squares of that column of X cannot be completely represented in one word of storage in most computers.

Rounding error is more likely to have a serious effect on the solution when the columns of X differ greatly in relative magnitude. If one column of X contains elements which are all small while another contains elements which are large the X'X matrix will contain a row in which one of the off diagonal elements is much larger than the diagonal element. In this case customary numerical procedures for inverting a matrix are such that the solution of the normal equations may be greatly affected by rounding error. This is so because these procedures yield results which are less affected by rounding error when the diagonal element is the largest element in each row of the matrix to be inverted.

All three conditions just discussed are usually present when fitting a polynomial equation. The X matrix for fitting a p^{th} order polynomial using n observations has the form

$$X = \begin{bmatrix} 1 & X_1 & X_1^2 & \cdots & X_1^p \\ 1 & X_2 & X_2^2 & \cdots & X_2^p \\ 1 & X_3 & X_3^2 & \cdots & X_3^p \\ \vdots & \vdots & \vdots & & \vdots \\ 1 & X_n & X_n^2 & \cdots & X_n^p \end{bmatrix}$$

The columns of X tend to be highly correlated in this situation therefore producing an X'X matrix whose determinant is small. In addition to ill-conditioning the last column contains elements which are the corresponding elements of the second column raised to the p^{th} power. Therefore each element of the last column usually has many significant digits and X'X may contain elements which cannot be precisely represented in a computer. Finally, unless the X_i's are very small the magnitudes of the elements of different columns will be greatly different thus producing an X'X matrix with several rows containing off diagonal elements which are larger than the diagonal element. For these reasons more caution than usual is recommended when fitting a polynomial equation.

References

The following publications are some of the many that describe available methods for solving equations when numerical difficulties are encountered.

Businger, Peter and Golub, H. [1965]. Linear Least Squares Solutions by Householder Transformations. Numerische Mathematik 7:269

Fox, L. [1964]. An Introduction to Numerical Linear Algebra. Claredon Press, Oxford.

Golub, G. H. [1965]. Numerical Methods for Solving Linear Least Squares Problems. Numerische Mathematik 7:206

Läuchli, P. [1961]. Jordan - Elimination and Ausgleichung nach Kleinsten Quadraten. Numerische Mathematik 3:226

Osborne, E. E. [1961]. On Least Squares Solutions of Linear Equations. Journal Assoc. Comput. Mach. 5:553

Riley, J. D. [1956]. Solving Systems of Linear Equations with a Positive Definite, Symmetric, but Possibly Ill-Conditioned Matrix. Math. Tables Aids Comput. 9:96

Siegel, I. H. [1961]. Least Squares "without normal equations". Trans. New York Acad. Sci. Ser. II 24:362