

COMPUTATION OF THE PSEUDOINVERSE OF

A MATRIX OF UNKNOWN RANK **/

by

Victor Pereyra^{*/} and J. B. Rosen

Abstract

A program is described which computes the pseudoinverse, and other related quantities, of an $m \times n$ matrix A of unknown rank. The program obtains least square solutions to singular and/or inconsistent linear systems $Ax = B$, where $m \leq n$ or $m > n$ and the rank of A may be less than $\min(m,n)$. A complete description of the program and its use is given, including computational experience on a variety of problems.

^{*/} On leave from Dept. of Mathematics and Computation Center, Buenos Aires University, Argentina.

**/

Reproduction in Whole or in part is Permitted for any Purpose of the United States Government. Prepared under NASA Grant Ns G 565 at Stanford University, Stanford, California.

I. Introduction.

A method for the computation of the pseudoinverse, and other related quantities, corresponding to an $m \times n$ matrix A of unknown rank r , has recently been described [5]. The method determines the pseudoinverse A^+ of A and a related matrix $A^\#$. The pseudoinverse has the property that given the linear system $Ax = b$, the solution $x_m = A^+b$ satisfies $\|Ax_m - b\| \leq \|Ax - b\|$ for all x , and $\|x_m\| \leq \|x\|$ for all x such that $\|Ax_m - b\| = \|Ax - b\|$. The minimum basic solution $x_b = A^\#b$ has the property that $\|Ax_b - b\| \leq \|Ax - b\|$ for all x , and x_b has at most r non-zero elements.

The computational difficulty for this problem arises primarily because the rank r is not known. In particular, it may be difficult to assign the correct rank if one or more of the singular values of A are small but non-zero [3]. Several other recent papers [1], [2], [4], on the computation of the pseudoinverse have not considered this important practical question.

The approach used here to handle this difficulty can be summarized as follows. The desired matrices $A^\#$ and A^+ are formed from a matrix B , which consists of linearly independent columns selected from A .

We would like to determine B so that it spans the same space as A , in which case B will contain r columns. Suppose we have a matrix B_q with q linearly independent columns selected from A , (where $q < r$) and the corresponding approximation A_q^+ to A^+ . Adding another linearly independent column of A to B_q , giving B_{q+1} , should give an improved approximation A_{q+1}^+ to the pseudoinverse. However, due to roundoff error in the calculation it may turn out for an ill-conditioned system that the new approximation is actually worse in the sense that $\|AA_{q+1}^+ - I\| > \|AA_q^+ - I\|$. Such a test is made in the pseudoinverse determination with the result that the effective rank of A (the number of columns in B) is the maximum possible consistent with minimizing the error $\|AA^+ - I\|$.

A closely related aspect of the method used here to compute the pseudoinverse is what might be called its "smoothing" property. In many practical situations one would like to obtain a solution to a linear system which is stable in the sense that small changes in the matrix elements do not cause large changes in the solution vector. In general, the solution $x = A^+b$, where A^+ is the true pseudoinverse, will not behave smoothly. In fact, the norm of x will increase without bound as a singular value of A approaches zero. This difficulty can be eliminated by imposing a predetermined upper bound on the norm of $(B'B)^{-1}$. This is accomplished by estimating the effect of adding a new column of A to B_q and only adding this new column to B_q if it does not cause any element of A_{q+1}^+ to exceed the bound. Details of this selection procedure and the manner in which it depends on the choice of the bound SUPER is discussed in the next section.

In Section 3 the use of the Algol 60 program, written to perform this algorithm, is described and suggested values for the input parameters are given. A large number of problems have been solved using an Extended Algol version of this program on a Burroughs B-5000 computer at Stanford University [7]. Several different kinds of tests have been performed:

- a) Very ill-conditioned matrices like the segments of the Hilbert matrix [6] have given a clear example of the smoothing property of the method.
- b) Random rectangular matrices of random sizes have been generated and the pseudoinverse have been computed. The sizes were allowed to vary between 1 and 25. In all the cases the results were satisfactory.
- c) Same as in b) but with random ranks. In every case the rank was correctly determined by the program.
- d) Random matrices of specified size covering a range of values of m and n were run in order to obtain time estimates for different size problems.
- e) A number of least-square problems, i.e., with $m \gg n$ and only one right-hand side.
- f) A variety of matrices for which an independent check on the accuracy of the solution was available.

Tests b) through f) showed that in reasonable problems in which the rank is well determined the program will work very well, while a) has shown that in very ill-conditioned cases the smoothing property of the method is effective.

These test results are discussed more fully in Section 4. The notation used in [5] will also be followed here.

Details of storage requirements are given in Section 5.

A copy of the program appears in the Appendix.

II. Program Description.

The method used to compute $A^\#$ and A^+ from B is essentially that given in Section 2 of [5]. For convenience we will repeat the key relations here. The pseudoinverse of the $m \times r$ matrix B of rank r is given by

$$(2.1) \quad B^+ = (B'B)^{-1} B'$$

The non-zero rows of the $n \times m$ matrix $A^\#$ then consists of the corresponding rows of B^+ . An $r \times n$ matrix of rank r is also obtained from B^+ according to

$$(2.2) \quad C = B^+A$$

Note that, if B contains all the independent columns of A , then $A = BC$. Finally, A^+ is obtained from C and B^+ by

$$(2.3) \quad A^+ = C'(CC')^{-1} B^+$$

If $r = n$ then the following equalities hold,

$$(2.4) \quad B = A, \quad A^\# = B^+ = A^+ \quad \text{and} \\ \|A^\#A - I\| = \|A^+A - I\| = 0$$

where $I = n \times n$ identity matrix.

Thus, it would appear to be unnecessary to compute A^+ from $A^\#$, since they are equal. Actually, because of round-off in the computation, $A^\#$ will not satisfy (2.4) exactly. We will prove now that it is worthwhile to compute A^+ from $A^\#$ by means of formulas (2.2) and (2.3), because in this fashion A^+ will be a more accurate representation of the pseudoinverse of A .

Lemma: Let A be a $m \times n$ matrix with rank n . Let \bar{A} be a $n \times m$ matrix with rank n , and such that,

$$\bar{A} A - I = C - I = E \quad \text{with} \quad \|E\| = \epsilon \leq 0.4$$

then the matrix $\bar{\bar{A}} = C'(CC')^{-1} \bar{A}$ satisfies

$$\|\bar{\bar{A}} A - I\| = O(\epsilon^2)$$

Proof: $\bar{\bar{A}} A = C'(CC')^{-1} C$

if C is nonsingular and symmetric then this identity gives trivially,

$$\bar{\bar{A}} A = I \quad .$$

However, in general, since $C = E + I$ we can write,

$$\bar{\bar{A}} A = (I + E')[(I + E)(I + E')]^{-1} (I + E)$$

also, since $(I + E)(I + E') = I + [E + E' + EE']$

we can write,

$$[(I + E)(I + E')]^{-1} = I - E - E' + O(E^2)$$

where $O(E^2)$ means that the remainder terms are at least quadratic in E, E' .

This expansion will converge if,

$$\|E + E' + EE'\| < 1 \quad .$$

But $\|E + E' + EE'\| \leq 2 \|E\| + \|E\|^2$

and the assumption $\|E\| \leq 0.4$ gives us the necessary bound.

Collecting terms we get the final result,

$$\bar{A} A = (I + E') [I - E - E' + O(E^2)] (I + E) = I + O(E^2)$$

or,

$$\|\bar{A} A - I\| = \|O(E^2)\| \leq K \cdot \epsilon^2$$

It is now clear that by putting the computed $A^\#$ instead of \bar{A} this Lemma insures us that A^+ (the \bar{A} of the Lemma) will give a better numerical approximation for the pseudoinverse of A .

This error squaring behavior was first observed experimentally, giving numerical evidence of the importance of performing the whole algorithm in this case.

The determination of B is based on the algorithm of Section 3 in [5], using the more sophisticated selection procedure described below.

We will now describe the program, an Algol 60 Procedure, called PSEUDOINVER, which will solve the matricial problem,

$$AX = \text{RHS}$$

where RHS is a $m \times t$ matrix.

a) The first part of PSEUDOINVER normalizes the matrix A by scaling each column so that its Euclidean norm is equal to one. The normalization constants are saved in order to get back to the original problem.

The search for independent columns of A is then made to determine the matrix B , according to the formulas described in section 3 of [5]. At this stage, the condition for a vector to be accepted as independent of the ones already included in the basis is that α , the square of the norm of the projection on the orthogonal subspace to that basis, be less than a quantity $ORTP$, which is an input parameter. Later we will discuss the appropriate choice of $ORTP$ and the other parameters appearing in the program.

As the columns chosen in this fashion might not necessarily be the first columns of A , a record is kept of the column number of the accepted vectors.

After all the columns have been inspected two situations can arise; either all the n columns of A have been accepted or some have been rejected. In the first case we have finished and the computations indicated at the beginning of this section are performed to get A^+ , $A^\#$, X_b and X_m . Other computed quantities are the residuals, $NXM = \|AX_m - RHS\|$, $NXB = \|AX_b - RHS\|$ corresponding to X_m and X_b , and $EST = \|BC - A\|$.

EST would be zero if the computation were performed exactly; in general EST will be very small for well-conditioned matrices and will increase with the ill-conditioning or if an almost dependent column is added to B .

If only $q < n$ columns of A are selected for inclusion in B then the basis thus constructed is called B_q and the second part of PSEUDOINVER is called.

b) The projections on the subspace orthogonal to B_q are computed for all the rejected columns. The Euclidean norm of each projection is computed,

(2.5)

$$\alpha_j = \|(I - B_q B_q^+) a_j\|^2$$

and the column corresponding to the maximum α_j (the most independent one) is stored in SAV.

c) A test is now performed which is based upon an estimation of the norm that $(B'_{q+1} B_{q+1})^{-1}$ would have if we were to include SAV in the basis. The norm used is $\|A\| = \max_i \sum_{j=1}^n |a_{ij}|$ and the estimate is derived from the formula,

$$(2.6) \quad (B'_{q+1} B_{q+1})^{-1} = \left(\begin{array}{c|c} (B'_q B_q)^{-1} & 0 \\ \hline 0 & 0 \end{array} \right) + \alpha_{q+1}^{-1} \begin{pmatrix} u_q \\ -1 \end{pmatrix} (u'_q \mid -1)$$

where $u_q = B_q^+ \text{SAV}$ and α_{q+1} is the square of the norm of the projection of SAV on the orthogonal subspace to that spanned by B_q . Then

$$\|(B'_{q+1} B_{q+1})\| \leq \text{ESTIM} = \|(B'_q B_q)^{-1}\| + \alpha_{q+1}^{-1} (1 + \sqrt{q})$$

If ESTIM is larger than SUPER (an input parameter) then SAV is rejected and B_q is taken as the final B.

This test avoids large elements in the pseudoinverse and gives the smoothing property discussed in the introduction.

d) If the test in c) is passed then the PROCEDURE GARBG, which computes all the matrices and quantities mentioned at the beginning of this section, is called and a second test is made. GARBG is used again, now with the basis B_q plus the column SAV. The test consists in comparing the values of $\|AX_m - \text{RHS}\|$, $\|AX_b - \text{RHS}\|$ and $\|BC - A\|$ obtained with one basis, with the corresponding ones obtained with the incremented basis. If all these values for B_{q+1} are smaller than for B_q then SAV is definitely

accepted. After shifting all the useful quantities, part (b) is repeated for the new basis B_{q+1} and so on, until either an exit is provided for one of the tests or the columns of A are exhausted. All the scalar products are performed in double precision. The block diagram in Fig. 1 shows the most essential parts of the program.

It is worth noting that this strategy has been dictated by the problem itself and achieves the best numerical pseudoinverse possible using the method of [5] and taking into account the numerical roundoff error of the computer being used. This strategy takes advantage of the step by step algorithm for determining B , and constructs an independent basis, the degree of independence being determined by the parameter ORTP. By picking the most independent vector among the remaining ones, and checking to see if this decreases the residuals (by taking this vector in the basis) we are answering in a direct manner the two questions: how many columns of A do we need to minimize the residual? and, among all the possible sets of independent columns, which set gives the best representation of the pseudoinverse?

III. Program Use.

As described in the previous section, several parameters are needed besides the matrix A . Now we will explain the use and possibilities of these parameters.

Input parameters.

m (integer) number of rows in A .

n (integer) number of columns in A .

t (integer) number of right-hand sides.

OPC (Boolean) If OPC is equal to 1 then the program will compute the matrices A^+ and $A^\#$, and the right-hand side $RHS = I(m \times m)$ will be automatically provided.

Moreover, OPC decides if in the test described in Section 2,d) the quality of the representation ($A = BC$) is controlled. That test is done only if $OPC = TRUE$. If OPC is equal to FALSE then RHS an $m \times t$ matrix has to be provided and the program will compute $X_B = A^\# \cdot RHS$, and $X_M = A^+ \cdot RHS$, matrices that will be printed out instead of A^+ and $A^\#$.

SUPER (real) It is the SUPER of Section 2,c). If an upper bound for the elements of A^+ is known then SUPER can be set to this bound to take advantage of the smoothing property of this method; otherwise it is suggested that 10^{14} be used. It should be noted that, in general, a larger value than 10^{14}

will increase computing time by throwing unnecessary decisions into the test of Section 2,d). On the other hand, much smaller values may completely eliminate from further consideration some columns which could be used to decrease the error.

ORTP (real) This parameter was described in Section 2, a). Small values for ORTP (around 10^{-4}) in general will accelerate the process because the first part (construction of a basis of strongly independent column) is the fastest and as many columns as possible should be accepted there. Nevertheless, there are at least two cases in which a more careful choice of ORTP may be important. If higher precision in the answers is desired (at the cost of increased computing time), then a larger value of ORTP should be used, say 0.05. This will allow the second part of the program to choose "better" columns.

The other delicate case occurs when the matrix is very ill-conditioned and the rank is therefore not well defined. Here the use of a relatively large ORTP is important. Again values around 0.05 are recommended.

Summarizing, in a reasonable, well behaved problem a recommended set of parameters is:

$$\text{SUPER} = 10^{14}, \quad \text{ORTP} = 10^{-4}.$$

If the representation becomes very bad ($\|A - BC\|$ too large), the problem is not well behaved and more burden should be passed on to the second and safer test by increasing SUPER and decreasing ORTP .

If the user has information that certain variables are more significant than others, this information can be used by ordering the matrix A so that the columns of A corresponding to these variables appear first. This will insure that these columns are considered first for inclusion in the basis B .

IV. Test Problems.

a) Square segments of the Hilbert matrix have been tried, sizes varying between 3 and 10.

For $5 \leq n \leq 10$ the rank found in each case was 4. The norm of the pseudoinverses remained below 10^3 while for the true pseudoinverse (the inverse in these cases) the norms ranged between 10^5 for $n=5$ and 10^{13} for $n=10$. The norm, $\|A - BC\|$ was around 10^{-5} for all cases.

As is well known, the ill-conditioning of the Hilbert matrix segments increase with their dimension. However, because of the smoothing property of the method a bounded and reasonably accurate representation for the pseudoinverse was always obtained.

b) Eighteen random matrices with random dimensions varying between 1 and 25 were generated and pseudoinverted. The norm $\|A - BC\|$ was always below 10^{-9} and the ranks were always found to be equal to $\min(m,n)$.

c) Given three random integers m, n and r in the interval $[1, 25]$ a routine generated two random matrices, $L (m \times r)$ and $R (r \times n)$. Multiplying them we obtained a matrix $A (m \times n)$ with rank at most equal to $r^{(1)}$. With 20 matrices generated in this way, the results were similar to b). In every case the rank r was correctly determined. For most of these cases the rank r was less than $\min(m,n)$, and of course was unknown for the program.

(1) This test was suggested by Professor Gene H. Golub.

d) For each pair of values (m,n) several random matrices were generated and pseudoinverted. Average values of $\|A-BC\|$ for these matrices with $m=10,20,30$ and $n=10,20,30$ are shown in Table I. For the same problems, average computation time on the Burroughs B5000 at Stanford Computation Center are shown in Table II.

In all these matrices the rank was the maximum possible, i.e., $\text{rank} = \min(m,n)$ and it was properly determined by the program.

TABLE I

m \ n	10	20	30
10	5×10^{-9}	9×10^{-9}	4.5×10^{-8}
20	3×10^{-10}	4×10^{-8}	2.9×10^{-7}
30	5.8×10^{-10}	6.9×10^{-9}	2.6×10^{-7}

$\|A - BC\|$

TABLE II

m \ n	10	20	30
10	7.6	13.9	20.7
20	16.5	49	91
30	27.3	92.5	180.6

Comp. time in seconds

For a 40×40 matrix the answers were:

$t = 413 \text{ sec.}$ $\|A - BC\| = 2.7 \times 10^{-7}$

$\text{rank} = 40$.

e) A common problem in many branches of applied sciences is the least squares fit, and is therefore one of the most important applications for this program. A related feature of the program is that, by ordering the variables, the user will be able to test their independence and eventually to decide if his model is appropriate to the phenomena being investigated. This is done by ordering

the matrix A so that the first coefficients correspond to the more important variables. The program will attempt to use these columns first to form the basis B . The necessity for such an ordering is clear from the fact that if we have n columns in A and the subspace spanned by these columns has dimension $(n-p)$ then we can construct with these columns as many as $\binom{n}{n-p}$ linearly independent sets.

In Figures II and III are shown the results obtained by running the program with least squares type problems. Again the elements of the matrices were generated randomly. Fig. II shows computation times on the B5000 for different values of m and $n=5, 10$. Fig. III shows the norm $\|A - BC\|$ for the same problems.

f) Matrices with exact known inverses were tried obtaining good results and accuracy. Of course this program should not be used to invert a matrix which is known to be nonsingular and well-conditioned, because it will be around four times slower than an efficient matrix inverter. The program has also been used to obtain the pseudoinverse of singular and almost singular matrices stemming from the discretization of integral equations of the first kind, and problems in pattern recognition.




FIGURE II




FIGURE III

V. Storage Requirements.

As all the array declarations are dynamical, the amount of storage depends on several parameters. If m , n , t are as before, and r is the final rank (number of accepted columns) then an estimate for the storage used in the PROCEDURE PSEUDOINVER is,

$$\text{Storage} = n^2 + 5mn + mr + 2nr + 2nt + \max(r^2, mn, mt)$$

the last term is present because in the PROCEDURE GARBG we have several independent blocks, and the storage corresponding to certain arrays is not simultaneously used.

If, as usual, r is not known, then it can be replaced by $\min(n,m)$.