SOLUTION OF A RANK-DEFICIENT LEAST SQUARES PROBLEM BY APPLYING ORTHOGONAL TRANSFORMATIONS

Keijo Inkilä
Jarmo Pirhonen

Helsinki University of Technology
Institute of Photogrammetry and Remote Sensing
Finland

Abstract

The paper deals with the determination of the rank of a matrix and the solution of a rank-deficient least squares problem. Several methods, which are based on the application of the elementary orthogonal transformations (Householder or Givens), are described. The methods are evaluated with respect to the numerical stability and computational efficiency. Special attention is paid on the exploitation of the sparsity of a matrix.

0. INTRODUCTION

The paper deals with the following problem:

> Given an m x n *matrix* A *(design matrix) with* $m > n$ *and rank* $(A) = k < n$ *and an m-vector* l *(observation vector) find an n-vector* $\hat{x}$ *such that*
>
> $$\hat{v}^T \hat{v} = (A\hat{x} - l)^T (A\hat{x} - l) = \text{minimum}. \qquad (1)$$

i.e. the solution of a rank-deficient least squares problem (lsd-problem).

In photogrammetry, a rank-deficient problem arises, e.g. if in point determination the datum is completely or partially undefined. Often we are only interested in the shape (and perhaps the size) of the object, and there is no need that the determined coordinates refer to a specific datum. This kind of situation is common particularly in close-range applications.

A rank defect also results easily in ls-approximation problems (DTM, image analysis) due to the devoid of data (uneven distribution of data). In particular, in piece-wise tensor product approximation it is frequently difficult to fulfill the condition for the full-rank problem.

The solutions of a lsd-problem presented in this paper are based on the use of orthogonal transformations, specifically on the use of

A revised version of the paper presented to the ISPRS Comm. III Symposium in Wuhan, 1990.

the QR-decomposition. A lsd-problem is essentially solved by solving successively two full-rank ls-problems. Therefore, we begin by introducing the QR-decomposition and by showing, how a full-rank ls-problem is solved by means of it.

1. FULL-RANK PROBLEM

If rank (A) = k = n, it can be shown that there exists an m x m orthogonal matrix, such that /4/

$$Q^T A = \begin{bmatrix} R \\ 0 \end{bmatrix}$$

or

$$Q \begin{bmatrix} R \\ 0 \end{bmatrix} = A \qquad \text{(QR-decomposition)}$$

where R is an n x n upper-triangular matrix. By defining

$$Q^T l = \begin{bmatrix} c \\ d \end{bmatrix} \begin{matrix} \}n \\ \}m-n \end{matrix}$$

and recalling that multiplication of a vector by an orthogonal matrix preserves the 2-norm (Euclidean norm), we obviously obtain the unique least squares solution $\hat{x}$ by solving the upper-triangular system

$$Rx = c \ .$$

It is easy to see that the residual sum of squares is

$$S = ||\hat{v}||^2 = ||A\hat{x}-l||^2 = \hat{v}^T \hat{v} = ||d||^2 = d^T d$$

and that the residual vector can also be computed as

$$\hat{v} = A\hat{x}-l = Q \begin{bmatrix} 0 \\ d \end{bmatrix} \ .$$

The reduction of the matrix A to the upper-triangular form is generally done by applying Householder or Givens transformations /4/. These elementary orthogonal transformations can be used to make one (Givens) or more (Householder) components of a vector to zero. The matrix Q obviously is the product of these elementary transformation matrices. Usually, one does not need to compute nor store the matrix Q explicitely.

It is important to note that the reduction of A to the upper triangular matrix R can be built up by processing the matrix A column by column (column-wise) or row by row (row-wise) (Fig.1) /4/. The first approach is well-known of its high numerical stability and high computational efficiency. On the other hand, processing row-wise makes recursive solutions particularly easy to compute and it is also advantageous considering the exploitation of sparsity.

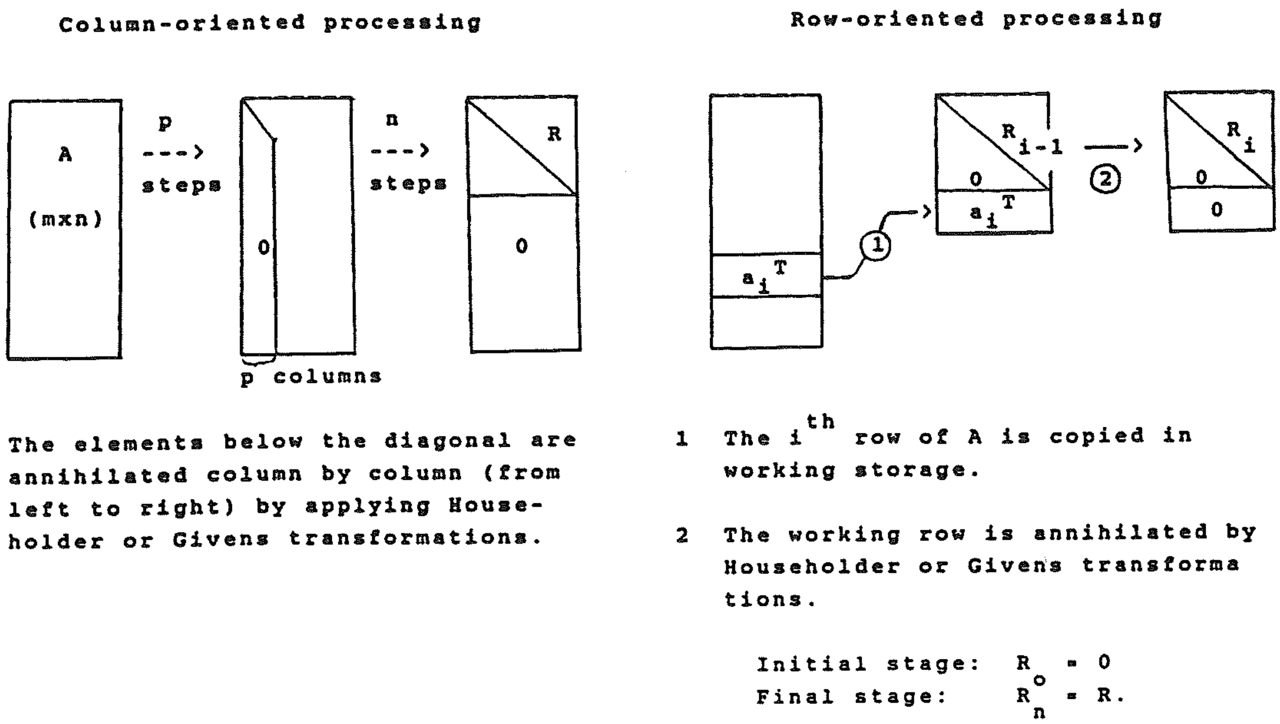

Figure 1. Column-oriented vs. row-oriented processing.

2. RANK-DEFICIENT PROBLEM

Let us assume now that

$$\text{rank}(A) = \text{rank}(A_1) = k < n$$

where $A = [\overset{k}{A_1} \; \overset{n-k}{A_2}]$, that is, only the first k columns of A are linearly independent and A is rank-deficient with the defect d=n-k. It can be shown that there exists an m x m-matrix Q such that

$$Q^T A = \begin{bmatrix} R_1 & R_2 \\ 0 & 0 \end{bmatrix} \begin{matrix} \} \; k \\ \} \; m-k \end{matrix}$$

where R_1 is a k x k upper-triangular matrix. Clearly, a unique solution is obtained only by imposing additional conditions. We consider in the following two particular solutions.

2.1 Basic solution

Let us define

$$Q^T l = g = \begin{bmatrix} g_1 \\ g_2 \end{bmatrix} \begin{matrix} \} \; k \\ \} \; m-k \end{matrix}$$

The basic solution $\bar{x}$ is simply obtained by solving the upper-triangular system $R_1 x_1 = g_1$ and defining

$$\bar{x} = \begin{bmatrix} x_1 \\ 0 \end{bmatrix} \begin{matrix} \} \; k \\ \} \; n-k \end{matrix}$$

A lsd-problem may obviously have as many as n!/k!(n-k)! basic solutions, but after deciding which components are set to zero the basic solution is unique.

The basic solution (also called reflexive ls-inverse solution or minimum rank ls-inverse solution) is often quite satisfactory. There are, however, strong reasons in favor of the minimum-norm solution to be presented in the next section.

2.2 Minimum-norm solution

The unique solution of minimum norm is obtained by requiring that in addition of the ls-criterion (1) the solution also fulfills the criterion

$$\tilde{x}^T \tilde{x} = \text{minimum}. \qquad (2)$$

The solution $\tilde{x}$ also is the pseudo-inverse solution, because the criterions (1) and (2) imply that both the observation space R^m and the parameter space R^n are equipped with the ordinary inner product.

Immediately an argument in favor of the minimum-norm solution is obvious: the condition (2) reduces unnecessary fluctuations of the components of the solution vector and thus restrains the appearance of numerical difficulties . But most of all, the solution $\tilde{x}$ also has a very appealing statistical property, namely that

$$\text{tr}\,(\Sigma_{\tilde{x}\tilde{x}}) = \text{minimum}$$

i.e. $\tilde{x}$ also is the minimum-variance estimate.

The minimum-norm solution is found perhaps most easily by means of vector space (geometrical) considerations. From Fig. 2 it is seen that $\tilde{x}$ is the orthogonal projection of $\bar{x}$ (any particular solution) to the row space $R(A^T)$ of A (recall that $R^n = R(A^T) \oplus N(A)$)

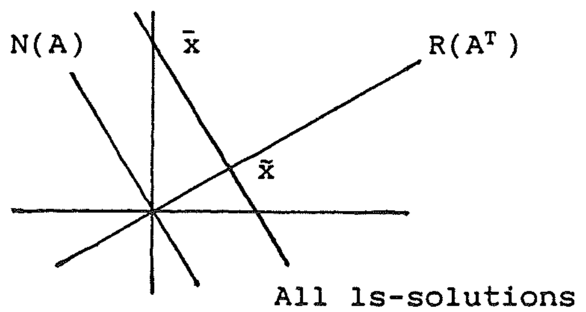

Figure 2. The basic solution vs. minimum-norm solution. ($R^n = R^2$, k = d = 1)

Obviously, given a basis of $R(A^T)$ or N(A) the minimum-norm solution is found by solving an additional full-rank ls-problem with the observation vector $\bar{x}$. Thus, two approaches can be distinguised depending on whether the solution is based on a basis of the row space of A, or the null space of A.

Row space methods

It is easy to establish that the columns of the n x k matrix

$$F = \begin{bmatrix} R_1^T \\ R_2^T \end{bmatrix}$$

constitute a basis of the row space $R(A^T)$. Thus,

$$\tilde{x} = F\hat{p}$$

where $\hat{p}$ is the solution of the full-rank ls-problem

$$Fp \simeq \bar{x} \qquad (3)$$

This problem can, of course, be solved by any method. If

$$F = WT = [W_1 \;\; W_2]\begin{bmatrix} T \\ 0 \end{bmatrix} \begin{matrix} \} \, k \\ \} \, n-k \end{matrix}$$

is the QR-decomposition of F, then after some manipulations we have

$$\tilde{x} = W_1 W_1^T \bar{x}$$

It can be shown further that $W_1^T \bar{x} = (T^T)^{-1} g_1$ and thus finally

$$\tilde{x} = W_1 (T^T)^{-1} g_1$$

The core of the solution is thus the computation of the complete orthogonal decomposition

$$Q^T A W = \begin{bmatrix} T^T & 0 \\ 0 & 0 \end{bmatrix}$$

Note finally that $R(W_1) = R(A^T)$, i.e. an orthogonal basis is constructed for the row space $R(A^T)$.

Null space methods

Due to the assumption that the columns of A_2 are linear combinations of the columns of A_1, we can solve a k x d matrix K from the equation

$$R_1 K = R_2$$

which satisfies the equation

$$A_1 K = A_2$$

or

$$[A_1 \;\; A_2]\begin{bmatrix} K \\ -I \end{bmatrix} = 0$$

or

$$AG = 0$$

Thus, the columns of the n x d matrix G constitute a basis of the null space N(A) and the minimum-norm solution $\tilde{x}$ is obtained as the residual vector of the full-rank ls-problem

$$Gq \simeq \bar{x} \qquad (4)$$

Consequently, if $H = [H_1 \;\; H_2]$ is an orthogonal n x n matrix of the QR-decomposition of G and if we define

$$H^T\bar{x} = \begin{bmatrix} H_1^T\bar{x} \\ H_2^T\bar{x} \end{bmatrix} = \begin{bmatrix} e_1 \\ e_2 \end{bmatrix} \begin{matrix} \} \; d=n-k \\ \} \; k \end{matrix}$$

then it is easy to derive the following explicit expressions for $\tilde{x}$:

$$\tilde{x} = H \begin{bmatrix} 0 \\ e_2 \end{bmatrix} = H_2 e_2$$

or

$$\tilde{x} = \bar{x} - H_1 e_1$$

Remarks:

(1) In general, d<<k and thus the computational effort required to solve the problem (4) is smaller than that of the problem (3).

(2) The application of the projection theorem yields

$$\tilde{x} = (I - G(G^TG)^{-1}G^T)\bar{x} = (I - H_1H_1^T)\bar{x} = S\bar{x}$$

Note that this is the S-transformation, when the solution is specifically constrained to lie in $R(A^T)$ /5/. Note also that the columns of H_1 are the eigenvectors of A^TA corresponding to the zero eigenvalues and that $R(H_1) = R(G) = N(A)$.

(3) If the rank defect is caused by the lacking datum information, then a basis of N(A) can also be found by simple geometrical considerations /1/.

3. STABILITY CONSIDERATIONS

We assumed above that $\text{rank}(A) = \text{rank}(A_1) = k$. Often these assumptions cannot, however, be made in practice. Then some additional measures are needed

(1) to guarantee that R_1 is non-singular (or $R(Q_1) = R(A)$) and possibly also
(2) to determine the rank (pseudorank) so that R_1 is reasonably well-conditioned

We discuss two approaches: the first one is applicable when A is processed column-wise, the second one when A is processed row-wise.

Column pivoting

The principle of column pivoting is very straightforward: columns are interchanged so that the unreduced column of largest norm is next to be reduced to the upper-triangular form. As a result of this ordering strategy the diagonal elements of R will be non-increasing in magnitude. In fact, they will satisfy the inequalities /4/

$$r_{kk}^2 \geq \sum_{i=k}^{j} r_{ij}^2 \qquad \text{for} \begin{cases} k = 1,\ldots,n-1 \\ j = k+1,\ldots,n \end{cases}$$

and

$$|r_{kk}| \geq |r_{ij}| \qquad \text{for} \quad i \geq k,\ j > k\ ,\ k = 1,\ldots,n-1.$$

Consequently, if $r_{k+1,k+1}$ is negligible (i.e. smaller than the pre-chosen tolerance value), then each column of the submatrix R_{22} of

$$R = \begin{bmatrix} R_{11} & R_{12} \\ 0 & R_{22} \end{bmatrix} \begin{matrix} \} \ k \\ \} \ n-k \end{matrix}$$

also has a negligible norm, which would suggest that the rank is k.

The inequalities above prove that the QR-decomposition with column pivoting is a very reliable method for the rank determination. Indeed, it can be regarded about as good as the singular value decomposition, which is generally regarded as the most reliable method /2/.

If this technique is properly applied to updating the norms, the additional computational amount for this pivoting technique is of order mn instead of mn^2 /2/.

Additional row processing

If the design matrix A is processed row-wise, the column pivoting cannot be applied. Then the stability of the computations can be improved by using the method proposed by Heath /3/.

The method is based on the testing of the diagonal elements of R after the completition of the row-wise processing. If $r_{jj} < \tau$ (τ is a pre-chosen tolerance value) for some j, then the following steps are taken:

1° The off-diagonal elements of the j^{th} row of R are transferred on the working row and the j^{th} row is filled by zeroes.
2° The working row is annihilated by normal processing.

After all diagonal elements have been tested (in increasing order), the rank of A is obtained as the number of the non-zero rows of R. The submatrices R_1 and R_2 are also easily extracted from this form.

It is worth emphasizing that the off-diagonal elements on the rows with small (negligible) diagonal elements are, in general, not negligible, but may be arbitrary large. They appear because of wrong placement of some rows of A into R during the row processing. The wrong decisions, in turn, are caused by inexact arithmetic (an element on the working row becomes non-zero due to errors).

The present method has been reported to work well in practice, although in theory it is somewhat less robust than the column pivoting method /3/. A very favourable feature of this approach comes out in the next chapter.

4. SPARSITY CONSIDERATIONS

A closer examination would reveal easily that in the QR-decomposition approach sparsity can be exploited to the same degree as in the normal equation approach in solving a full-rank problem /3/. It should be noted specifically that when the processing is done row-wise, the intermediate fill-in is confined to the working row only allowing use of a static data structure for R.

If additional row processing is used for the enhancement of numerical stability, the column ordering of A can be chosen by considering the exploitation of sparsity only and sparsity can be exploited efficiently in the solution of a lsd-problem, too (note that d=n-k is usually small implying that the problem (4) is a small dense problem).

5. COMPUTATION OF COFACTOR MATRICES

For statistical purposes it is often required to compute the cofactor matrix for a minimum-norm solution. One approach is to apply the propagation law of random errors to the expression $\tilde{x} = S\bar{x}$, where $S = (I - H_1 H_1^T)$ yielding

$$Q_{\tilde{x}\tilde{x}} = S Q_{\bar{x}\bar{x}} S^T$$

$Q_{\tilde{x}\tilde{x}}$ can be computed efficiently and stably as

$$Q_{\tilde{x}\tilde{x}} = CC^T$$

where C is n x k matrix which is solved from the triangular system

$$R_1^T C^T = S_1^T$$

and S_1 is a n x k submatrix of $S = [S_1 \;\; S_2]$. Compared to the full-rank case, additional work thus is due to the computation of S and due to the fact that C is now a rectangular matrix, not a triangular matrix. Note that because S is symmetric its computation requires only $d(n^2+n)/2$ multiplications and, in addition, d is usually small.

Alternatively, the cofactor matrix $Q_{\tilde{x}\tilde{x}}$ can be obtained by defining

$$S = H_2 H_2^T$$

or

$$S = W_1 W_1^T$$

Note that the computation of S requires now $k(n^2+n)/2$ multiplications and usually k>>d.

6. AN EXAMPLE

To demonstrate the difference between the basic solution and the minimum-norm solution, let us consider the following surface fitting problem:

- The data points are distributed as shown in Fig. 3. Note that in addition to a gap in the middle of the area there are less points on the right-hand border.
- A bicubic tensor-product spline is used as an approximating function. The knot lines of the spline are shown by solid lines in Fig. 3.

A closer inspection would reveal that the resulting ls-problem is rank-deficient (the redundancy is 184 and the rank-defect is 4).

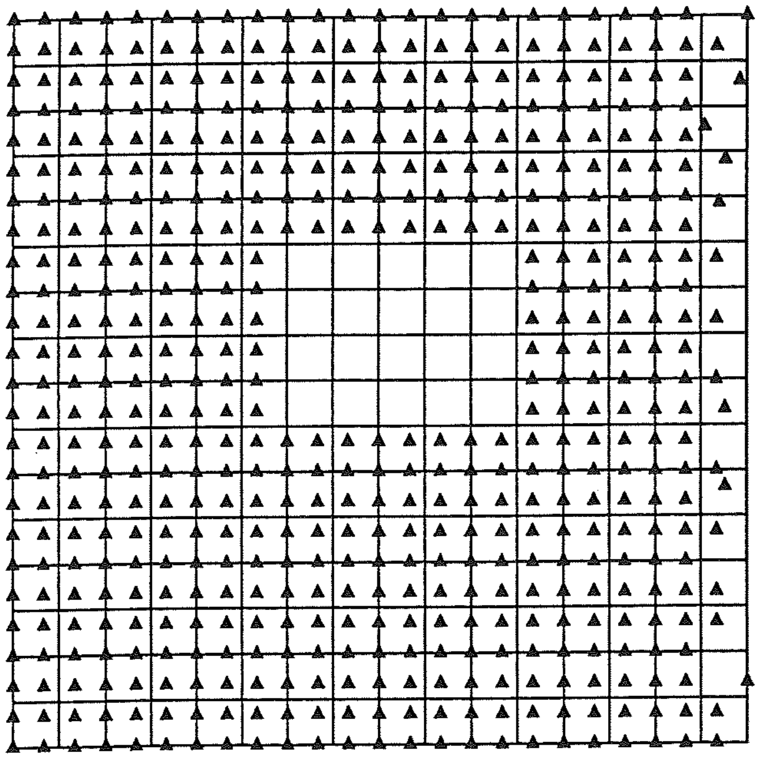

Figure 3. The data point distribution and the knot lines.

If exactly the same value is applied for each observation (underlying function is a horizontal plane), then the approximating surfaces shown in Figures 4 and 5 are obtained by using the basic solution and the minimum-norm solution, respectively.

In accordance with the expectations fluctuations are remarkably reduced, when the unique solution is based on the minimum-norm condition.

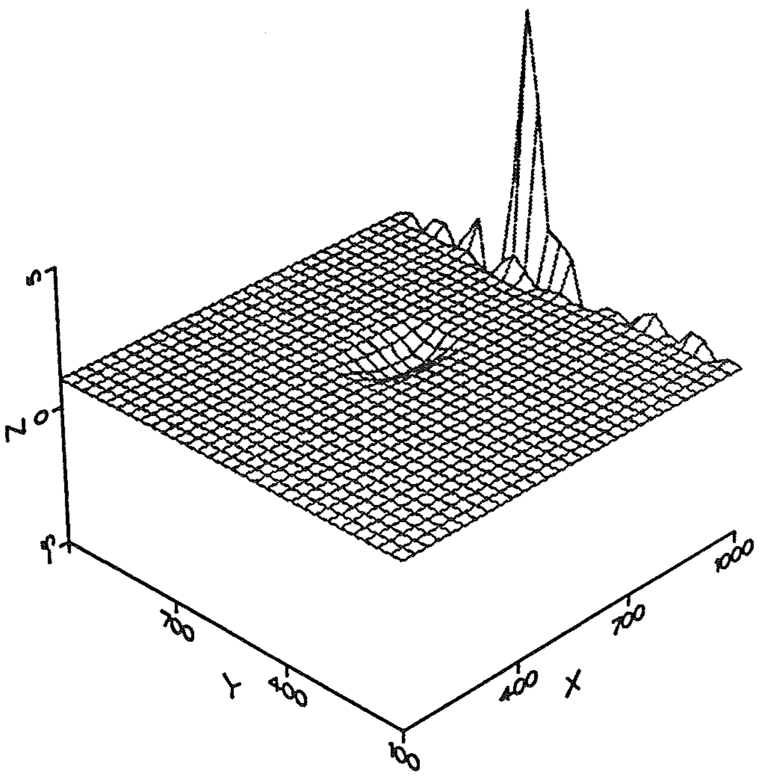

Figure 4. The approximating surface obtained by applying the basic solution.

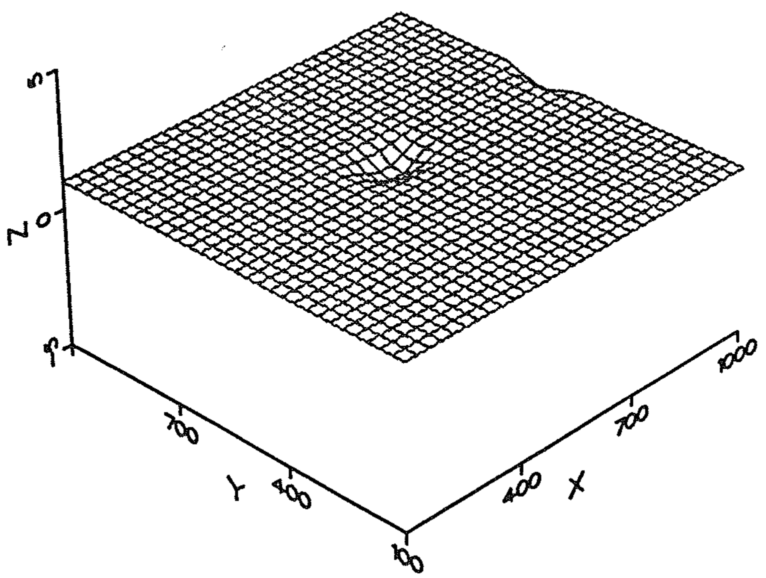

Figure 5. The approximating surface obtained by applying the minimum-norm solution.

7. REFERENCES

/1/ Fraser, C.S.: Optimization of Precision in Close-Range Photogrammetry. Photogrammetric Engineering and Remote Sensing, Vol. 48, No. 4, April 1984.

/2/ Golub, G.H., van Loan, C.F.: Matrix Computations. North Oxford Academic, Oxford, 1983.

/3/ Heath, M.T.: Some Extensions of an Algorithm for Sparse Linear Least Squares Problems. SIAM J. Sci. Stat. Comput., Vol. 3, No. 2, June 1982.

/4/ Lawson, C.L., Hanson, R.J.: Solving Least Squares Problems. Prentice-Hall, New Jersey, 1974.

/5/ Molenaar, M.: S-transformations and Artificial Covariance Matrices in Photogrammetry. ITC-Journal 1981-1, 1981.