

Euclidean distances and least squares problems for a given set of vectors [☆]

Xiao-Wen Chang ^{*}, Christopher C. Paige

School of Computer Science, McGill University, Montreal, Quebec, Canada, H3A 2A7

Available online 9 February 2007

Abstract

Given an $m \times n$ matrix A , n Euclidean distances, those from each column to the space spanned by the remaining columns of A , are considered. An elegant relationship between A , these Euclidean distances, and the solutions of n simple linear least squares problems arising from A is derived. When A has full column rank, from this a useful expression for these Euclidean distances is immediately obtained. The theory is then used to greatly improve the efficiency of an algorithm used in communications.

© 2007 IMACS. Published by Elsevier B.V. All rights reserved.

Keywords: Least squares; Residuals; Euclidean distance; QR factorization; Permutations; Communications

1. Introduction

Given n real m -vectors, n simple linear least squares (LS) problems can be formed, each using one of these n vectors as the right-hand side of a LS problem, while the remaining $n - 1$ vectors form the coefficient matrix of the problem. The 2-norm of each least squares residual vector is the Euclidean distance from the chosen right-hand side vector to the space spanned by the other $n - 1$ vectors.

This paper shows an elegant relationship between these Euclidean distances (2-norms of the LS residuals) and the LS solutions, and gives a general formula for each Euclidean distance. Then an application in communications is considered, in which the QR factorization of a certain column permutation of an $m \times n$ matrix is needed. Based on the formulae for the Euclidean distances, a greatly improved algorithm is proposed for finding the permutations and obtaining the QR factorization.

In the paper bold upper case letters and bold lower case letters are used to denote matrices and vectors respectively. The identity matrix is denoted by I and its i th column by e_i . Sometimes MATLAB notation is used to denote a submatrix. Specifically, if A is a matrix, then $A(:, j)$ is the j th column of A , $A(:, j_1 : j_2)$ denotes the submatrix formed by columns j_1 to j_2 , and $A(i_1 : i_2, j_1 : j_2)$ the submatrix formed by rows i_1 to i_2 and columns j_1 to j_2 .

[☆] This research was supported by NSERC of Canada Grants RGPIN217191-03 for Xiao-Wen Chang and RGPIN9236-01 for Christopher C. Paige.

^{*} Corresponding author.

E-mail addresses: chang@cs.mcgill.ca (X.-W. Chang), paige@cs.mcgill.ca (C.C. Paige).

URLs: <http://www.cs.mcgill.ca/~chang> (X.-W. Chang), <http://www.cs.mcgill.ca/~chris> (C.C. Paige).

2. The main result

The following theorem gives some interesting and useful general results.

Theorem 2.1. Given an $m \times n$ real matrix $A = [a_1, \dots, a_n]$, let $D(A) \equiv \text{diag}(\delta_1(A), \dots, \delta_n(A))$ where for $i = 1, \dots, n$, $\delta_i(A)$ is the Euclidean distance from a_i to the space spanned by all other columns of A , i.e.,

$$\delta_i(A) = \min_{x_i} \|a_i - [a_1, \dots, a_{i-1}, a_{i+1}, \dots, a_n]x_i\|_2 = \|A\bar{x}_i\|_2.$$

Here if $\hat{x}_i \equiv [\hat{x}_{1,i}, \dots, \hat{x}_{i-1,i}, \hat{x}_{i+1,i}, \dots, \hat{x}_{n,i}]^T$ is any vector solving the above LS problem (\hat{x}_i might not be unique), we define

$$\bar{x}_i \equiv [\hat{x}_{1,i}, \dots, \hat{x}_{i-1,i}, -1, \hat{x}_{i+1,i}, \dots, \hat{x}_{n,i}]^T, \quad \bar{X} \equiv [\bar{x}_1, \bar{x}_2, \dots, \bar{x}_n]. \tag{1}$$

Then there is a relationship between A , the “solution” vectors \bar{x}_i , and the Euclidean distances $\delta_i(A)$:

$$A^T A \bar{X} = -D^2. \tag{2}$$

When A has full column rank, the diagonal elements of the symmetric matrix $\bar{X}D^{-2} = -(A^T A)^{-1}$ show that

$$\delta_i(A)^2 = \frac{1}{e_i^T (A^T A)^{-1} e_i}, \quad i = 1, 2, \dots, n, \tag{3}$$

while for $i, j = 1, 2, \dots, n$, with $i \neq j$, the off-diagonal elements give

$$\frac{\hat{x}_{j,i}}{\delta_i(A)^2} = \frac{\hat{x}_{i,j}}{\delta_j(A)^2} = -e_i^T (A^T A)^{-1} e_j. \tag{4}$$

This relates the least squares residual norms (or Euclidean distances) $\delta_i(A)$ and $\delta_j(A)$, to the least squares solutions \hat{x}_i and \hat{x}_j . Combining (3) and (4) relates the least squares solutions \hat{x}_i and \hat{x}_j to the elements of $(A^T A)^{-1}$:

$$\hat{x}_{j,i} e_i^T (A^T A)^{-1} e_i = \hat{x}_{i,j} e_j^T (A^T A)^{-1} e_j = -e_i^T (A^T A)^{-1} e_j. \tag{5}$$

Proof. With the definitions in (1), for $i = 1, \dots, n$ denote

$$A_i \equiv [a_1, \dots, a_{i-1}, a_{i+1}, \dots, a_n], \quad d_i \equiv a_i - A_i \hat{x}_i = -A \bar{x}_i. \tag{6}$$

The normal equations for the LS problem $\min_{x_i} \|a_i - A_i x_i\|_2$ give

$$A_i^T d_i = 0. \tag{7}$$

Thus from (6) and (7) it follows that

$$a_i^T d_i = (d_i + A_i \hat{x}_i)^T d_i = d_i^T d_i = \delta_i(A)^2. \tag{8}$$

Combining (7) and (8), we obtain

$$A^T d_i = e_i \delta_i(A)^2,$$

which, by the expression for d_i in (6), is equivalent to

$$A^T A \bar{x}_i = -e_i \delta_i(A)^2.$$

But this gives (2); and then (3), (4) and (5) follow for full column rank A . \square

Corollary 2.1. Let full column rank A have the QR factorization

$$A = Q \begin{bmatrix} R \\ \mathbf{0} \end{bmatrix}, \tag{9}$$

where Q is an $m \times m$ orthogonal matrix and R is an $n \times n$ nonsingular upper triangular matrix with elements r_{ij} . Then

$$r_{kk}^2 = \delta_k(A(:, 1:k))^2 = \frac{1}{e_k^T (A(:, 1:k))^T A(:, 1:k))^{-1} e_k}, \quad k = 1, \dots, n. \tag{10}$$

Proof. From the upper triangular structure of \mathbf{R} it can be seen that $|r_{kk}|$ is the Euclidean distance from $\mathbf{R}(:, k)$ to the space spanned by the columns of $\mathbf{R}(:, 1 : k - 1)$, or equivalently the Euclidean distance from \mathbf{a}_k to the space spanned by $\mathbf{a}_1, \dots, \mathbf{a}_{k-1}$. Thus the first equality in (10) holds. The second equality in (10) follows from (3) in Theorem 2.1. \square

Remark 2.1. The formula (3) in Theorem 2.1 can be proved in a slightly more straightforward way. Suppose we interchange the i th column and the n th column of full column rank \mathbf{A} and find the corresponding QR factorization:

$$\mathbf{A}\mathbf{P} = \mathbf{Q} \begin{bmatrix} \mathbf{R} \\ \mathbf{0} \end{bmatrix}, \quad \mathbf{P} \equiv [\mathbf{e}_1, \dots, \mathbf{e}_{i-1}, \mathbf{e}_n, \mathbf{e}_{i+1}, \dots, \mathbf{e}_{n-1}, \mathbf{e}_i].$$

Then

$$\mathbf{P}^T (\mathbf{A}^T \mathbf{A})^{-1} \mathbf{P} = \mathbf{R}^{-1} \mathbf{R}^{-T}.$$

Comparing the (n, n) elements of both sides, we obtain

$$\mathbf{e}_n^T \mathbf{P}^T (\mathbf{A}^T \mathbf{A})^{-1} \mathbf{P} \mathbf{e}_n = 1/r_{nn}^2.$$

But from the upper triangular structure of \mathbf{R} , we see that $|r_{nn}|$ is the Euclidean distance from the last column of $\mathbf{A}\mathbf{P}$ (which is \mathbf{a}_i) to the space spanned by the other columns of $\mathbf{A}\mathbf{P}$, so $|r_{nn}| = \delta_i(\mathbf{A})$. Thus

$$\delta_i(\mathbf{A})^2 = 1/\mathbf{e}_i^T (\mathbf{A}^T \mathbf{A})^{-1} \mathbf{e}_i.$$

Remark 2.2. The relation between $|r_{kk}|$ and $\delta_k(\mathbf{A}(:, 1 : k))$ given by the first equality in (10) is known, see [2, p. 104].

3. An application

The theory given in Section 2 might have several uses. Here we show one use for (3).

In solving integer linear least squares problems which may arise in decoding or detection in communications, one often computes the QR factorization of \mathbf{A} with column pivoting in order to make the later computation (a search process) efficient (see, e.g., [1] and [3]). The permutation matrix \mathbf{P} is usually determined such that the $|r_{kk}|$ are large for large k and small for small k (note that $|r_{11}r_{22} \cdots r_{nn}|$ is fixed). In [3], the so-called vertical Bell Labs layered space-time (V-BLAST) optical detection ordering given in [4] was proposed for the column permutations. This permutation strategy determines the columns of the permuted matrix $\mathbf{A}\mathbf{P}$ from the last to the first. Let \mathcal{J}_k denote the set of column indices for the not yet chosen columns when the k th column of the permuted matrix $\mathbf{A}\mathbf{P}$ is to be determined (where $k = n, n - 1, \dots, 1$). Then this strategy chooses the $p(k)$ th column of the original matrix \mathbf{A} to be the k th column of the permuted matrix $\mathbf{A}\mathbf{P}$ that we seek, as follows:

$$p(k) = \arg \max_{j \in \mathcal{J}_k} \mathbf{a}_j^T [\mathbf{I} - \mathbf{A}_{k,j} (\mathbf{A}_{k,j}^T \mathbf{A}_{k,j})^{-1} \mathbf{A}_{k,j}^T] \mathbf{a}_j, \quad (11)$$

where $\mathbf{A}_{k,j}$ is the $m \times (k - 1)$ matrix formed by the columns \mathbf{a}_i with $i \in \mathcal{J}_k - \{j\}$. We can easily show that $\mathbf{a}_j^T [\mathbf{I} - \mathbf{A}_{k,j} (\mathbf{A}_{k,j}^T \mathbf{A}_{k,j})^{-1} \mathbf{A}_{k,j}^T] \mathbf{a}_j$ in (11) is the squared Euclidean distance from \mathbf{a}_j to the space spanned by the columns of $\mathbf{A}_{k,j}$. But Corollary 2.1 tells us that $|r_{kk}|$ is the Euclidean distance from the k th column of $\mathbf{A}\mathbf{P}$ to the space spanned by the first $k - 1$ columns of $\mathbf{A}\mathbf{P}$. Thus for $k = n, n - 1, \dots, 1$, $\mathbf{a}_{p(k)}$ is the column which makes $|r_{kk}|$ maximum among all the not yet chosen columns when we determine the k th column of the permuted matrix $\mathbf{A}\mathbf{P}$, i.e., the resulting $|r_{kk}|$ is the largest which can be found at this step.

The following is an obvious way to find this permutation matrix \mathbf{P} . We first compute the QR factorization of \mathbf{A} as in (9) and record $|r_{nn}|$. In order to determine the last column of the permuted \mathbf{A} , we interchange the last column of \mathbf{R} with its j th column for $j = 1, 2, \dots, n - 1$. After each interchange, we compute the QR factorization of the permuted \mathbf{R} , which can be done in an efficient way (taking $O(n^2)$ flops) by using the structure of the permuted \mathbf{R} , and record the corresponding $|r_{nn}|$. By comparing the n possible values of $|r_{nn}|$, we determine the desired last column of the permuted \mathbf{A} , and keep the updated \mathbf{R} -factor for use in the next step. We see that the cost of these computations is $O(n^3)$ flops. Then in order to determine the desired second to last column of the permuted \mathbf{A} , we apply a similar process to the $(n - 1) \times (n - 1)$ submatrix $\mathbf{R}(1 : n - 1, 1 : n - 1)$ of the updated \mathbf{R} we have just obtained — to obtain a new updated upper triangular $\mathbf{R}(1 : n - 1, 1 : n - 1)$ having the largest $|r_{n-1, n-1}|$. After repeating the above process $n - 1$ times, all the permutations are determined.

For the final \mathbf{R} we hope that $0 < |r_{11}| \leq \dots \leq |r_{nn}|$, but this is not necessarily true as the example $\mathbf{A} = \begin{bmatrix} 9 & 6 \\ 0 & 8 \end{bmatrix}$ shows, since exchanging the columns leads to $|r_{22}| = 7.2$, showing that \mathbf{A} is already in the chosen final order.

The above algorithm is numerically stable, but it needs $O(n^4)$ flops. Later a faster algorithm is given. For motivation consider the QR factorization of any permutation of the columns of \mathbf{A} , and its first k columns,

$$\mathbf{A}\mathbf{P} = \mathbf{Q} \begin{bmatrix} \mathbf{R} \\ \mathbf{0} \end{bmatrix}, \quad \mathbf{A}\mathbf{P}(:, 1:k) = \mathbf{Q}(:, 1:k)\mathbf{R}(1:k, 1:k). \tag{12}$$

From this and Corollary 2.1 we see that if \mathbf{P} is the permutation matrix arising from V-BLAST, then for $k = n, n - 1, \dots, 1$,

$$|r_{kk}| = \delta_k(\mathbf{A}\mathbf{P}(:, 1:k)) = \max_{1 \leq j \leq k} \delta_j(\mathbf{A}\mathbf{P}(:, 1:k)),$$

or equivalently, using (3), the aim of V-BLAST is to find \mathbf{P} such that

$$\mathbf{e}_k^T [\mathbf{P}(:, 1:k)^T \mathbf{A}^T \mathbf{A} \mathbf{P}(:, 1:k)]^{-1} \mathbf{e}_k = \min_{1 \leq j \leq k} \mathbf{e}_j^T [\mathbf{P}(:, 1:k)^T \mathbf{A}^T \mathbf{A} \mathbf{P}(:, 1:k)]^{-1} \mathbf{e}_j. \tag{13}$$

For the faster algorithm, first compute the QR factorization of \mathbf{A} and denote this

$$\mathbf{A} = \tilde{\mathbf{Q}} \begin{bmatrix} \tilde{\mathbf{R}} \\ \mathbf{0} \end{bmatrix}. \tag{14}$$

Next compute $\tilde{\mathbf{L}} \equiv \tilde{\mathbf{R}}^{-T}$, which takes about $n^3/3$ flops (cf. [2, p. 119]). Note for the diagonal elements, $\tilde{l}_{ii}\tilde{r}_{ii} = 1$, and from (3)

$$\delta_n(\mathbf{A})^{-2} = \mathbf{e}_n^T (\mathbf{A}^T \mathbf{A})^{-1} \mathbf{e}_n = \mathbf{e}_n^T \tilde{\mathbf{R}}^{-1} \tilde{\mathbf{R}}^{-T} \mathbf{e}_n = \mathbf{e}_n^T \tilde{\mathbf{L}}^T \tilde{\mathbf{L}} \mathbf{e}_n = \tilde{l}_{nn}^2 = \tilde{r}_{nn}^{-2},$$

so if $\mathbf{L} = \mathbf{R}^{-T}$ in (12), to maximize $\delta_n(\mathbf{A}\mathbf{P})$ over all permutations \mathbf{P} we can either maximize r_{nn}^2 , or minimize l_{nn}^2 (and so on for further steps). To do this, we compute a QL factorization of $\tilde{\mathbf{L}}$ with a particular form of pivoting:

$$\tilde{\mathbf{L}}\tilde{\mathbf{P}} = \tilde{\mathbf{Q}}\tilde{\mathbf{L}}. \tag{15}$$

This is done using permutations and orthogonal rotations, and will cost no more than about n^3 flops (note that this does not count the cost of explicitly forming $\tilde{\mathbf{Q}}$, since it is not necessary in the application). We just work on $\tilde{\mathbf{L}}$ until finally it becomes $\tilde{\mathbf{L}}$. In the first step we move the column of $\tilde{\mathbf{L}}$ with smallest 2-norm (suppose it is the p th column) to become the last column of $\tilde{\mathbf{L}}\tilde{\mathbf{P}}$ (moving each of columns $p + 1, \dots, n$ down one). We then apply $n - p$ rotations from the left of $\tilde{\mathbf{L}}\tilde{\mathbf{P}}$ to bring it back to lower triangular form (cf. [2, p. 134]), which we call ‘the updated $\tilde{\mathbf{L}}$ ’. In general, for $k = n, n - 1, \dots, 2$, in the $(n - k + 1)$ th step of the reduction process we determine the index p such that column p of the updated submatrix $\tilde{\mathbf{L}}(1:k, 1:k)$ has the *smallest* 2-norm among all of its columns. Then we move column p to be the k th column of the updated $\tilde{\mathbf{L}}$. After that we apply $k - p$ orthogonal rotations to zero $\tilde{\mathbf{L}}(1:k - 1, k)$. This step costs about $3(k - 2p + 2)k$ flops, from which it can be concluded that the total cost of the algorithm to obtain the QL factorization (15) is at most about n^3 flops (depending on the permutations required). Thus we have for the column norms of $\tilde{\mathbf{L}}(1:k, 1:k)$

$$\mathbf{e}_k^T \tilde{\mathbf{L}}(1:k, 1:k)^T \tilde{\mathbf{L}}(1:k, 1:k) \mathbf{e}_k = \min_{1 \leq j \leq k} \mathbf{e}_j^T \tilde{\mathbf{L}}(1:k, 1:k)^T \tilde{\mathbf{L}}(1:k, 1:k) \mathbf{e}_j. \tag{16}$$

But since $\tilde{\mathbf{L}} = \tilde{\mathbf{R}}^{-T}$, from (15) and (14) it follows that

$$\tilde{\mathbf{L}}^{-1} \tilde{\mathbf{L}}^{-T} = \tilde{\mathbf{P}}^T \tilde{\mathbf{L}}^{-1} \tilde{\mathbf{L}}^{-T} \tilde{\mathbf{P}} = \tilde{\mathbf{P}}^T \tilde{\mathbf{R}}^T \tilde{\mathbf{R}} \tilde{\mathbf{P}} = \tilde{\mathbf{P}}^T \mathbf{A}^T \mathbf{A} \tilde{\mathbf{P}}.$$

Comparing the $k \times k$ leading principal submatrices of both sides, we have

$$\tilde{\mathbf{L}}(1:k, 1:k)^{-1} \tilde{\mathbf{L}}(1:k, 1:k)^{-T} = \tilde{\mathbf{P}}(:, 1:k)^T \mathbf{A}^T \mathbf{A} \tilde{\mathbf{P}}(:, 1:k).$$

Therefore

$$\tilde{\mathbf{L}}(1:k, 1:k)^T \tilde{\mathbf{L}}(1:k, 1:k) = [\tilde{\mathbf{P}}(:, 1:k)^T \mathbf{A}^T \mathbf{A} \tilde{\mathbf{P}}(:, 1:k)]^{-1}.$$

From this and (16) it follows that for $k = n, \dots, 1$

$$\mathbf{e}_k^T [\tilde{\mathbf{P}}(:, 1:k)^T \mathbf{A}^T \mathbf{A} \tilde{\mathbf{P}}(:, 1:k)]^{-1} \mathbf{e}_k = \min_{1 \leq j \leq k} \mathbf{e}_j^T [\tilde{\mathbf{P}}(:, 1:k)^T \mathbf{A}^T \mathbf{A} \tilde{\mathbf{P}}(:, 1:k)]^{-1} \mathbf{e}_j, \tag{17}$$

which with (13) shows that $\bar{\mathbf{P}}$ in (15) has the desired properties of \mathbf{P} in V-BLAST, and so we have found the permutations far more efficiently.

If \mathbf{A} is close to rank deficient, then the numerical computation $\bar{\mathbf{L}} \equiv \bar{\mathbf{R}}^{-T}$ might occasionally result in $\bar{\mathbf{P}}$ in (15) being not quite correct. That is, it might not give (17) for every k . Such an occurrence (hopefully rare in practice) could alter the number of steps in the later search process in solving the integer least squares problem, but would have no other deleterious effects.

So far we have found $\bar{\mathbf{P}}$, whereas our ultimate aim was to compute the QR factorization of $\mathbf{A}\bar{\mathbf{P}}$. There are two obvious ways to do this. From (14), $\bar{\mathbf{R}} = \bar{\mathbf{L}}^{-T}$, and (15), we obtain

$$\mathbf{A}\bar{\mathbf{P}} = \bar{\mathbf{Q}} \begin{bmatrix} \bar{\mathbf{L}}^{-T} \bar{\mathbf{P}} \\ \mathbf{0} \end{bmatrix} = \bar{\mathbf{Q}} \begin{bmatrix} \tilde{\mathbf{Q}} \\ \mathbf{I} \end{bmatrix} \begin{bmatrix} \tilde{\mathbf{L}}^{-T} \\ \mathbf{0} \end{bmatrix},$$

which, with $\mathbf{R} \equiv \tilde{\mathbf{L}}^{-T}$, gives the QR factorization we have sought:

$$\mathbf{A}\bar{\mathbf{P}} = \mathbf{Q} \begin{bmatrix} \mathbf{R} \\ \mathbf{0} \end{bmatrix}, \quad \mathbf{Q} \equiv \bar{\mathbf{Q}} \begin{bmatrix} \tilde{\mathbf{Q}} \\ \mathbf{I} \end{bmatrix}. \quad (18)$$

Computing $\mathbf{R} \equiv \tilde{\mathbf{L}}^{-T}$ again takes about $n^3/3$ flops. The problem is that \mathbf{R} may lose accuracy unnecessarily since it is obtained via the computation of the inverses of two triangular matrices. The other approach is to use the permutation matrix $\bar{\mathbf{P}}$ obtained in (15) and compute the QR factorization $\bar{\mathbf{R}}\bar{\mathbf{P}} = \hat{\mathbf{Q}}\hat{\mathbf{R}}$, so that from (14)

$$\mathbf{A}\bar{\mathbf{P}} = \bar{\mathbf{Q}} \begin{bmatrix} \bar{\mathbf{R}}\bar{\mathbf{P}} \\ \mathbf{0} \end{bmatrix} = \bar{\mathbf{Q}} \begin{bmatrix} \hat{\mathbf{Q}}\hat{\mathbf{R}} \\ \mathbf{0} \end{bmatrix} = \mathbf{Q} \begin{bmatrix} \mathbf{R} \\ \mathbf{0} \end{bmatrix}, \quad \mathbf{Q} \equiv \bar{\mathbf{Q}} \begin{bmatrix} \hat{\mathbf{Q}} \\ \mathbf{I} \end{bmatrix}, \quad \mathbf{R} \equiv \hat{\mathbf{R}},$$

which is a continuation of the original QR factorization (14), and so is also numerically stable. This approach takes at most about n^3 flops (depending on the permutations), which at worst could be a little more costly than the first approach, but it is more numerically reliable and so is more preferable in general.

Acknowledgements

The authors would like to thank Gene Golub and Michael Saunders for their helpful comments. The first author is indebted to them for their hospitality when he was on sabbatical at Stanford University where part of this research work was done.

References

- [1] E. Agrell, T. Eriksson, A. Vardy, K. Zeger, Closest point search in lattices, *IEEE Trans. Inform. Theory* 48 (2002) 2201–2214.
- [2] A. Björck, *Numerical Methods for Least Squares Problems*, SIAM, Philadelphia, PA, 1996.
- [3] M.O. Damen, H. El Gamal, G. Caire, On maximum-likelihood detection and the search for the closest lattice point, *IEEE Trans. Inform. Theory* 49 (2003) 2389–2402.
- [4] G.J. Foschini, G.D. Golden, R.A. Valenzuela, P.W. Wolniansky, Simplified processing for high spectral efficiency wireless communication employing multi-element arrays, *IEEE J. Select. Areas Commun.* 17 (1999) 1841–1852.