COMPUTATION OF THE C-S DECOMPOSITION, WITH APPLICATION TO SIGNAL PROCESSING

Charles Van Loan
Department of Computer Science
Cornell University
Ithaca, NY 14853

Jeffrey Speiser
Naval Ocean Systems Center
San Diego, CA 92152

Abstract

The CS decomposition and the Generalized Singular Value Decomposition are closely related. Indeed, the stable computation of the latter requires calculation of the former. In this paper we show how a number of important numerical activities in signal and image processing benefit from exploitation of these decompositions.

Introduction

Algorithmic developments in numerical linear algebra during the last twenty years have led to a number of "slogans" that are now widely appreciated in the scientific computing community. Here are some examples:

(1) You do not need to form A^TA to solve the least square problem $\min \|Ax - b\|_2$. (Compute the QR decomposition.)

(2) You do not need to diagonalize A^TA to obtain the eigenvalues of A^TA. (Compute the singular value decomposition.)

(3) You do not need to form $B^{-1}A$ to obtain the generalized eigenvalues of $A - \lambda B$. (Compute the QZ decomposition.)

The QR, singular value, and QZ decompositions are discussed in Golub and Van Loan (1983). So also is a little known decomposition called the "CS decomposition". The point of this paper is to generate some enthusiasm for the CS decomposition and in particular, the following "slogan":

(4) You do not need to simultaneously diagonalize A^TA and B^TB in order to solve the generalized singular problem $A^TAx = \lambda B^TBx$. (Compute the CS decomposition associated with A and B's Q factor.)

As we will show, there are several important computational problems in signal and image processing to which this fourth dictum applies. These applications alone suggest that the CS decomposition deserves a heightened profile. But to convince the reader further of the centrality of the CS decomposition, we show how it can assist in the development of a powerful parallel eigenvalue algorithm and how it provides yet another perspective of the FFT.

The CS Decomposition

Suppose that

$$Q = \begin{bmatrix} Q_{11} & Q_{12} \\ Q_{21} & Q_{22} \end{bmatrix} \begin{matrix} m_1 \\ m_2 \end{matrix}$$

$$\begin{matrix} n_1 & n_2 \end{matrix}$$

satisfies $Q^HQ = I$. Without loss of generality we assume that each Q_{ij} has at least as many rows as columns. The gist of the CS decomposition is that the SVD's of the Q_{ij} are highly related. In particular, there exist unitary U_1, U_2, V_1, and V_2 so that

$$\begin{bmatrix} U_1 & 0 \\ 0 & U_2 \end{bmatrix}^H \begin{bmatrix} Q_{11} & Q_{12} \\ Q_{21} & Q_{22} \end{bmatrix} \begin{bmatrix} V_1 & 0 \\ 0 & V_2 \end{bmatrix} = \begin{bmatrix} \mathrm{diag}(c_1,\ldots,c_{n_1}) & \mathrm{diag}(s_1,\ldots,s_{n_2}) \\ -\mathrm{diag}(s_1,\ldots,s_{n_1}) & \mathrm{diag}(c_1,\ldots,c_{n_2}) \end{bmatrix} \qquad (1)$$

Here, $c_i^2 + s_i^2 = 1$ for $i = 1:\max\{n_1,n_2\}$. Since the singular values of the Q_{ij} are cosines and sines, the decomposition (1) is called the "CS" decomposition".

The theoretical aspects of the CS decomposition are covered in Davis and Kahan(1970), Stewart(1977), Paige and Saunders(1981), and Van Loan (1984). Practical algorithms are specified in Stewart (1983) and Van Loan (1985). The methods presented in these papers deal with some subtle numerical difficulties that prevent one from approaching the CS decomposition as "just a bunch of calls to the LINPACK SVD routine."

The $A^HAx = \lambda B^HBx$ Problem

Suppose $A \in C^{m_1 xn}$ $(m_1 \geq n)$ and $B \in C^{m_2 xn}$ $(m_2 \geq n)$. Set $S_1 = A^HA$ and $S_2 = B^HB$. Since the S_i are Hermitian semidefinite, there exists a nonsingular X so

$$\begin{aligned} X^H(A^HA)X &= D_A = \mathrm{diag}(\alpha_1,\ldots,\alpha_n) \qquad \alpha_i \geq 0 \\ X^H(B^HB)X &= D_B = \mathrm{diag}(\beta_1,\ldots,\beta_n) \qquad \beta_i \geq 0 \end{aligned} \tag{2}$$

See Golub and Van Loan (1983,p313ff) for details. Note that if $X = [\, x_1,\ldots, x_n]$ is a column partitioning, then

$$\begin{aligned} A^HAx_i &= X^{-H}D_Ae_i = \alpha_i X^{-H}e_i \\ B^HBx_i &= X^{-H}D_Be_i = \beta_i X^{-H}e_i \end{aligned} \tag{3}$$

where e_i is the i-th column of the n-by-n identity. Thus, if $\lambda_i = \alpha_i/\beta_i$ is defined we have

$$A^HAx_i = \lambda_i B^HBx_i \tag{4}$$

and so we speak of (λ_i, x_i) as an eigenpair for $A^HA - \lambda B^HB$.

In some applications an orthonormal basis for

$$S(i_1,\ldots,i_p) = \mathrm{span}\,\{x_{i_1},\ldots, x_{i_p}\} \tag{5}$$

is required where $\{\lambda_{i_1},\ldots, \lambda_{i_p}\}$ is some designated subset of generalized eigenvalues. For example, one may require an orthonormal basis for the space associated with the three largest eigenvalues.

We mention that the generalized eigenvalue problem $A^HA - \lambda B^HB$ is also known as the generalized singular value problem. The square roots of the λ_i are the generalized singular values of the pair (A,B).

The CS Decomposition and the $A^HAx = \lambda B^HBx$ Problem

Suppose

$$\begin{bmatrix} A \\ B \end{bmatrix} = \begin{bmatrix} Q_1 \\ Q_2 \end{bmatrix} R \qquad \begin{aligned} \mathrm{size}(A) &= \mathrm{size}(Q_1) \\ \mathrm{size}(B) &= \mathrm{size}(Q_2) \end{aligned} \tag{6}$$

is the QR factorization of "A-over-B". Since $Q_1^HQ_1 + Q_2^HQ_2 = I$ it follows from the first block column of (1) that there exist unitary U_1, U_2, and V so

$$\begin{aligned} U_1^HQ_1V &= C = \mathrm{diag}(c_1,\ldots,c_n) \\ U_2^HQ_2V &= S = \mathrm{diag}(s_1,\ldots,s_n) \end{aligned} \tag{7}$$

When combined with (6) we have

$$A = U_1 C (V^H R)$$
$$B = U_2 S (V^H R) \qquad (8)$$

Thus, if we set $X = (V^H R)^{-1}$, $D_A = C^T C$, and $D_B = S^T S$, then $X^H(A\ A)X = D_A$ and $X^H(B\ B)X = D_B$. It follows that if $R^{-1}V = [x_1, \ldots, x_n]$ is a column partitioning and $\lambda_i = (c_i/s_i)^2$, then (4) holds.

Tacit in the above derivation is the assumption that R is nonsingular. This is guaranteed if the null spaces of A and B intersect only trivially. We assume this to be the case. (It is possible to proceed when A and B have both zero a common nonzero vector, but this doesn't add to the discussion.)

Eigenvector Beamforming

In the MUSIC algorithm of Schmidt (1981) it is necessary to compute an orthonormal basis associated with the minimum generalized singular value of the pair (A,B). This basis defines a "noise subspace" which is subsequently used to identify arrival directions. The following procedure was proposed in Speiser and Van Loan (1984):

Algorithm SVL

1. Compute the QR factorization

$$\begin{bmatrix} A \\ B \end{bmatrix} \quad \begin{bmatrix} Q_1 \\ Q_2 \end{bmatrix} R \qquad \begin{matrix} A,\ Q_1 \in C^{m_1 xn} \\ B,\ Q_2 \in C^{m_2 xn} \end{matrix}$$

2. Compute the SVD

$$U_2^H Q_2 V = \mathrm{diag}(s_1, \ldots, s_n) \qquad s_1 \geq s_2 \geq \cdots \geq s_n$$

and determine r so that

$$s_1 \geq \cdots \geq s_r \geq s_1(1 - \varepsilon) > s_{r+1} \geq \cdots \geq s_n$$

where ε is a small tolerance that depends on the accuracy of the data and/or the machine precision. Note that c_1/s_1 is the smallest generalized singular value and so r is a computed estimate of its multiplicity.

3. If $V^H R = [\, y_1, \ldots, y_n \,]$ is a column partitioning and Z is a unitary matrix chosen so

$$[y_1, \ldots, y_r]\, Z = T$$

is upper triangular, then the first r columns of Z define the required subspace.

This procedure is perfectly reliable. It is unaffected by singularity in $A^H A$ and $B^H B$ and does not require the formation of cross product matrices. The full CS decomposition of Q_1 and Q_2 is not even needed--just an ordinary SVD of Q_2. (This is numerically "safe" because we are computing the large singular values of Q_2. See Van Loan (1985) for more details.)

The Hotelling Trace Condition

The Hotelling Trace Condition (HTC) is one of several methods in image processing that can be used to measure separability of a set of object classes. An "object" is an N-vector and may consist, for example, of the values that define a 2-dimensional intensity distribution. An "object class" is a collection of sample objects that share some common characteristics, e.g., all x-ray photographs that reveal a certain type of crystal configuration.

Suppose f is a sample object and that we have r object classes. The covariance matrix for the k-th class is given by

$$C_k = E[\, (f - m_k)(f - m_k)^T \,] \qquad E = \text{expected value}$$

where m_k is the ensemble average over the k-th class. The average covariance of the object classes is given

by

$$S_2 = \sum p_k C_k \qquad \text{(N-by-N)}$$

where $p_1,\ldots,p_r$ are a priori probabilities. (p_k = the probability of class k occurring.) The interclass covariance matrix is given by

$$S_1 = \sum p_k E[\,(m_k - m)(m_k - m)^T\,] \qquad \text{(N-by-N)}$$

and relates the scattering of the object class means m_k to the overall expectation

$$m = \sum_{k=1}^{r} p_k m_k \quad .$$

The HTC is defined to be the scalar

$$J(S_1,S_2) = \text{trace}(S_2^{-1} S_1) \quad .$$

When this is small the object classes are similar while a large J value implies that the object classes are more distinct, the desirable situation. See Smith et al (1985) for details.

Note that if we solve the matrix equation

$$S_2 X = S_1$$

then

$$J(S_1,S_2) = \sum_{1}^{N} x_{ii} = \sum_{1}^{N} \lambda_i \qquad \lambda_i = \text{e.v. of } S_1 - \lambda S_2$$

However, if S_2 is ill-conditioned then J can be an unreliable measure of separability. A better procedure is to recognize that the trace is the sum of the eigenvalues of $S_1 - \lambda S_2$. Since the S_i have to be estimated in practice, e.g.,

$$S_1 = A^T A \qquad S_2 = B^T B \qquad (9)$$

we see that

$$J(S_1,S_2) \cong \sum (c_i/s_i)^2$$

where the c_i and s_i are the singular values of Q_1 and Q_2 in

$$\begin{bmatrix} A \\ B \end{bmatrix} = \begin{bmatrix} Q_1 \\ Q_2 \end{bmatrix} R \qquad Q_1^T Q_1 + Q_2^T Q_2 = I$$

In other words, the HTC can be estimated by computing the CS decomposition of A-over-B's Q factors.

Feature Selection Using the Hotelling Trace Condition

Picking a set of L features that maximally preserve class separability with respect to the HTC entails finding a matrix H (L-by-N) so that

$$J(HS_1H^T\ ,\ HS_2H^T)$$

is maximized. As is shown in Fukunga(1972), this can be done by setting the i-th row of H to be x_i^T where

$$S_1 x_i = \lambda_i S_2 x_i \qquad \lambda_1 \geq \lambda \geq \cdots \geq \lambda_N$$

Actually, it is not necessary for the rows of H to be made up of the L largest eigenvectors. Rather, the range of H^T need only span the space defined by these L vectors. Again, S_1 and S_2 come to us in the form(9) and so a slight modification of the SVL algorithm is applicable:

(1) Compute the QR factorization

$$\begin{bmatrix} A \\ B \end{bmatrix} = \begin{bmatrix} Q_1 \\ Q_2 \end{bmatrix} R$$

(2) Compute the CS decomposition

$$Q_1 = U_1 C V^T \qquad C = \mathrm{diag}(c_i), \quad c_1 \geq c_2 \geq \cdots \geq c_N$$

$$Q_2 = U_2 S V^T \qquad S = \mathrm{diag}(s_i)$$

(3) Set $V^T R = [\, y_1, \ldots, y_N]$ and compute orthogonal H so $[y_1, \ldots, y_L]H = T$ is upper triangular.

It should be recognized that the matrices A and B are very large in practice and special techniques need to be used.

The Block Jacobi Procedure for Eigenvalues

Suppose that $n = kp$ and that

$$A = \begin{bmatrix} A_{11} & \cdots & A_{1p} \\ \vdots & & \vdots \\ A_{p1} & \cdots & A_{pp} \end{bmatrix} \qquad A_{ij} \quad R^{kxk}$$

is symmetric. In the block Jacobi procedure for diagonalizing A one repeatedly solves subproblems of the form

$$\begin{bmatrix} Q_{11} & Q_{12} \\ Q_{21} & Q_{22} \end{bmatrix}^T \begin{bmatrix} A_{ii} & A_{ij} \\ A_{ji} & A_{jj} \end{bmatrix} \begin{bmatrix} Q_{11} & Q_{12} \\ Q_{21} & Q_{22} \end{bmatrix} = \begin{bmatrix} D_1 & 0 \\ 0 & D_2 \end{bmatrix} \tag{10}$$

where the Q matrix is orthogonal. In the scalar case ($k = 1$, $p = n$) Q has the form

$$\begin{bmatrix} Q_{11} & Q_{12} \\ Q_{21} & Q_{22} \end{bmatrix} = \begin{bmatrix} \cos(\theta) & \sin(\theta) \\ -\sin(\theta) & \cos(\theta) \end{bmatrix}$$

There are two choices for θ and by choosing the smaller of the two angles, convergence is much more rapid. (Otherwise the oscillation of the diagonal entries retards the rate at which the off-diagonal entries become small.)

Our experience with the block algorithm indicates that it is crucial to solve the subproblems correctly. Otherwise, block Jacobi applied to A as a p-by-p block matrix requires the same number of sweeps that the scalar algorithm would require by regarding A as an n-by-n matrix. This is very troublesome since in the parallel setting it is advantageous to implement the algorithm with the largest possible block size. (p is usually twice the number of available processors.) To get true block performance it would be desireable for the c-s values associated with Q to satisfy $|c_i| \geq |s_i|$, $i = 1{:}k$. A heuristic way to achieve this is to sort the eigenvalues in each subproblem (10). By doing this we find that the number of block sweeps is a function of p and not n. The sorting has a way of forcing the matrices Q_{12} and Q_{21} to zero as the iteration progresses. See Van Loan (1985) and Bischof and Van Loan (1985) for further details.

The CS Decomposition and the Fast Fourier Transform

Our last topic is just a curiosity but it helps to underscore the CS decomposition's fundamental nature. Suppose $n = 2^k$ and $x \in C^n$. The discrete Fourier transform of x is given by

$$y = F(n)x \quad , \quad F(n) = (f_{pq}) \quad , \quad f_{pq} = [\cos(2\pi pq/n) - i \sin(2\pi pq/n)]/\sqrt{n}$$

Here, the columns of F(n) are indexed from 0 to n-1. Ordinarily, such a matrix vector product requires n^2 flops. However, because of the special structure of F(n), it is possible to compute the transform y in just nlog(n) time. The "gist" of F(n)'s special structure has to do with the CS decomposition of the matrix

$$G(n) = F(n)P(n)$$

where P(n) is the bit reverse permutation. The action of P(n) is best illustrated with an example:

$$[z_0\ z_1\ z_2\ z_3\ z_4\ z_5\ z_6\ z_7\ ]P(8) = [\ z_0\ z_4\ z_2\ z_6\ z_1\ z_5\ z_3\ z_7\ ]$$

The order of the subscripts in the result is obtained by reversing the bits in natural sequence, 3 = 011 becomes 110 = 5.

Now it can be shown that the unitary matrix G(n) has the following CS decomposition:

$$G(n) = \frac{1}{\sqrt{2}} \begin{bmatrix} I & -I \\ I & I \end{bmatrix} \cdot \begin{bmatrix} G(n/2) & 0 \\ 0 & -D(n/2)G(n/2) \end{bmatrix}$$

where we define the diagonal matrices $D(m) \in C^{m \times m}$ by $D(m) = \mathrm{diag}(1, \omega, \ldots, \omega^{k-1})$ where $\omega = \exp(2\pi i/m)$. Note that if we have the transforms G(n/2)w and G(n/2)z then

$$G(n) \begin{bmatrix} w \\ z \end{bmatrix}$$

can be obtained in O(n) flops. Recurring on this idea implies that $y = F(n)x = G(n)[P(n)^T x]$ can be obtained with O(nlog(n)) work.

The point in all this is that the "V factors" in G(n)'s CS decomposition involve G(n/2) and the "U factors are identities.

Acknowledgements

The research in this paper was supported through ONR. The first author's contract is N-00014-83-K-0640.

References

C.Bischof and C.Van Loan (1986), "Computing the Singular Value Decomposition on a ring of array processors", Proceedings of the Workshop on Large Scale Eigenvalue Problems, North Holland.

C. Davis and W. Kahan (1970), "The rotation of eigenvectors by a perturbation III", SIAM J. Numer. Anal., 7, 1-46.

K. Fukunaga (1972), Introduction to Statistical Pattern Recognition , Academic Press, New York.

G.H. Golub and C. Van Loan (1983), Matrix Computations , Johns Hopkins University Press, Baltimore, Md.

C.C. Paige and M.A. Saunders (1981), "Towards a generalized singular value decomposition", SIAM J. Numer. 18, 398-405.

R.O.Schmidt (1981), "A signal subspace approach to multiple emitter location and spectral estimation", PhD dissertation, Department of Electrical Engineering, Stanford University.

W.E. Smith, H. Barrett, B.D. Fiete, K.J. Myers, and T.J. Roney (1985), "Ascheme for the optimization of imaging systems based on object-class separability", Acta Polytechnica Scandinavica Physics Series, Vol. 149, 254-257.

J. Speiser and C. Van Loan (1984), "Signal processing computations using the generalized singular value decomposition", Real Time Signal Processing VII, Proc. SPIE, Vol. 485, paper 495-32, pp. 47-55.

G.W. Stewart (1977), "On perturbation of pseudo-inverses, projections, and linear least squares problems", SIAM Review 19, 634-662.

G.W. Stewart (1983), "An algorithm for computing the CS decomposition of a partitioned orthogonal matrix", Numer. Math. 40, 297-306.

C. Van Loan (1985), "Computing the CS and Generalized Singular Value Decompositions", Numer. Math. 46, 479-91.

C. Van Loan (1986), "The block Jacobi method for computing the singular value decomposition," in Proceedings of MTNS-85, 7th International symposium on the Mathematical Theory of Networks and Systems, C. Byrnes and A. Lindquist (eds).