

Credit:

UTokyo Online Education, Special Lecture 2017 Gilbert Strang

License:

You may use this material in a page unit for educational purposes. Except where otherwise noted, this material is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives license in a page unit.

<http://creativecommons.org/licenses/by-nc-nd/4.0/>

These lecture materials contain images, etc. which the University of Tokyo has obtained permission for use from third parties and images, etc. provided under each license. These images, etc. may not be used separately from the lecture materials. When using separate images, etc., please comply with the terms of use of each right holder.



Multiplying and Factoring Matrices

Gilbert Strang
MIT
gs@math.mit.edu

I believe that the right way to understand matrix multiplication is **columns times rows** :

$$AB = \begin{bmatrix} \mathbf{a}_1 & \dots & \mathbf{a}_n \end{bmatrix} \begin{bmatrix} \mathbf{b}_1^T \\ \vdots \\ \mathbf{b}_n^T \end{bmatrix} = \mathbf{a}_1 \mathbf{b}_1^T + \dots + \mathbf{a}_n \mathbf{b}_n^T. \quad (1)$$

Each column \mathbf{a}_k of an m by n matrix multiplies a row of an n by p matrix. The product $\mathbf{a}_k \mathbf{b}_k^T$ is an m by p matrix of *rank one*. The sum of those rank one matrices is AB .

All columns of $\mathbf{a}_k \mathbf{b}_k^T$ are multiples of \mathbf{a}_k , all rows are multiples of \mathbf{b}_k^T . The i, j entry of this rank one matrix is $a_{ik} b_{kj}$. The sum over k produces the i, j entry of AB in “the old way.” Computing with numbers, I still find AB by **rows times columns** (inner products)!

The central ideas of matrix analysis are perfectly expressed as matrix factorizations :

$$\boxed{A = LU \quad A = QR \quad S = Q\Lambda Q^T \quad A = X\Lambda Y^T \quad A = U\Sigma V^T}$$

The last three, with eigenvalues in Λ and singular values in Σ , are often seen as column-row multiplications (a sum of outer products). The spectral theorem for S is a perfect example. The first two are Gaussian elimination (LU) and Gram-Schmidt orthogonalization (QR). We aim to show that those are also clearly described using rank one matrices.

The spectral theorem $S = Q\Lambda Q^T$ A real symmetric matrix S is diagonalized by its orthonormal eigenvectors. The eigenvalues λ_i enter the diagonal matrix Λ . They multiply the eigenvectors \mathbf{q}_i in the columns of Q . Then $\lambda_i \mathbf{q}_i$ is a column of $Q\Lambda$. The column-row multiplication $(Q\Lambda)Q^T$ has the familiar form

$$S = \lambda_1 \mathbf{q}_1 \mathbf{q}_1^T + \dots + \lambda_n \mathbf{q}_n \mathbf{q}_n^T. \quad (2)$$

To see that S times \mathbf{q}_j produces $\lambda_j \mathbf{q}_j$, multiply every term $\lambda_i \mathbf{q}_i \mathbf{q}_i^T$ by \mathbf{q}_j . By orthogonality, the only surviving term has $i = j$. That term is $\lambda_j \mathbf{q}_j$ because $\mathbf{q}_j^T \mathbf{q}_j = 1$.

Of course the proof of the spectral theorem requires construction of the \mathbf{q}_j .

Elimination $A = LU$ is the result of Gaussian elimination in the usual order, starting with an invertible matrix A and ending with an upper triangular U . The key idea is that the matrix L linking U to A contains the *multipliers* — the numbers ℓ_{ij} that multiply row j when it is subtracted from row $i > j$ to produce $U_{ij} = 0$.

The “magic” is that those separate steps do not interfere, when we undo elimination and bring U back to A . The numbers ℓ_{ij} fall into place in L —but that key fact can take patience to verify in a classroom. Here we look for a different approach. The column-times-row idea makes the steps of elimination transparently clear.

Step 1 of elimination Row 1 of U is row 1 of A . Column 1 of L is column 1 of A , divided by the first pivot a_{11} (so that $\ell_{11} = 1$). Then the product $\ell_1 \mathbf{u}_1^T$ extends the first row and column of A to a rank-one matrix. So the difference is a matrix A_2 of size $n - 1$ bordered by zeros in row 1 and column 1:

$$\text{Step 1} \quad A = \ell_1 \mathbf{u}_1^T + \begin{bmatrix} 0 & \mathbf{0}^T \\ \mathbf{0} & A_2 \end{bmatrix}. \quad (3)$$

Step 2 acts in the same way on A_2 . The row \mathbf{u}_2^T and the column ℓ_2 will start with a single zero. The essential point is that elimination on column 1 removed the matrix $\ell_1 \mathbf{u}_1^T$ from A .

$$A = \begin{bmatrix} 2 & 3 \\ 4 & 11 \end{bmatrix} \quad \mathbf{u}_1^T = [2 \quad 3] \quad \ell_1 = \begin{bmatrix} 1 \\ 2 \end{bmatrix} \quad \ell_1 \mathbf{u}_1^T = \begin{bmatrix} 2 & 3 \\ 4 & 6 \end{bmatrix}$$

$$A - \ell_1 \mathbf{u}_1^T = \begin{bmatrix} 0 & 0 \\ 0 & 5 \end{bmatrix} = \ell_2 \mathbf{u}_2^T = \begin{bmatrix} 0 \\ 1 \end{bmatrix} [0 \quad 5]. \quad \text{So } LU = \begin{bmatrix} 1 & 0 \\ 2 & 1 \end{bmatrix} \begin{bmatrix} 2 & 3 \\ 0 & 5 \end{bmatrix}.$$

When elimination reaches A_k , there are $k - 1$ zeros at the start of each row and column. Those zeros in \mathbf{u}_k and ℓ_k produce an upper triangular matrix U and a lower triangular L . The diagonal entries of U are the pivots (not zero). The diagonal entries of L are all 1's.

The linear system $A\mathbf{x} = \mathbf{b}$ is reduced to two triangular systems governed by L and U :

$$\text{Solve } L\mathbf{c} = \mathbf{b} \quad \text{and solve } U\mathbf{x} = \mathbf{c}. \quad \text{Then } A\mathbf{x} = LU\mathbf{x} = L\mathbf{c} = \mathbf{b}.$$

Forward elimination leaves $U\mathbf{x} = \mathbf{c}$, and back-substitution produces \mathbf{x} . To assure nonzero pivots, this LU decomposition requires every leading square submatrix of A (from its first k rows and columns) to be invertible.

Gram-Schmidt orthogonalization $A = QR$ The algorithm combines independent vectors $\mathbf{a}_1, \dots, \mathbf{a}_n$ to produce orthonormal vectors $\mathbf{q}_1, \dots, \mathbf{q}_n$. Subtract from \mathbf{a}_2 its component in the direction of \mathbf{a}_1 . Normalize at each step to unit vectors \mathbf{q} :

$$\mathbf{q}_1 = \frac{\mathbf{a}_1}{\|\mathbf{a}_1\|} = \frac{\mathbf{a}_1}{r_{11}} \quad \mathbf{q}_2 = \frac{\mathbf{a}_2 - (\mathbf{q}_1^T \mathbf{a}_2) \mathbf{q}_1}{\|\mathbf{a}_2 - (\mathbf{q}_1^T \mathbf{a}_2) \mathbf{q}_1\|} = \frac{\mathbf{a}_2 - r_{12} \mathbf{q}_1}{r_{22}}.$$

As with elimination, this is clearer when we recover the original vectors \mathbf{a}_1 and \mathbf{a}_2 from the final \mathbf{q}_1 and \mathbf{q}_2 :

$$\mathbf{a}_1 = r_{11} \mathbf{q}_1 \quad \mathbf{a}_2 = r_{12} \mathbf{q}_1 + r_{22} \mathbf{q}_2. \quad (4)$$

In this order we see why R is triangular. At each step, \mathbf{q}_1 to \mathbf{q}_k span the same subspace as \mathbf{a}_1 to \mathbf{a}_k . We can establish the Gram-Schmidt factorization $A = QR = \mathbf{q}_1 r_1^T + \dots + \mathbf{q}_n r_n^T$ as follows:

The first column \mathbf{q}_1 is the first column \mathbf{a}_1 divided by its length r_{11}

The first row r_1^T contains the inner products $\mathbf{q}_1^T \mathbf{a}_k$.

Subtracting the rank one matrix $\mathbf{q}_1 \mathbf{r}_1^T$ leaves a matrix A_2 whose columns are **all orthogonal to \mathbf{q}_1** :

$$A = \mathbf{q}_1 \mathbf{r}_1^T + \begin{bmatrix} \mathbf{0} & A_2 \end{bmatrix}. \quad (5)$$

This is the analog of equation (3) for elimination. There we had a row of zeros above A_2 . Here we have columns of A_2 orthogonal to \mathbf{q}_1 . In two lines, this example reaches equation (5) :

$$A = \begin{bmatrix} \mathbf{a}_1 & \mathbf{a}_2 \end{bmatrix} = \begin{bmatrix} 6 & 2 \\ 8 & 6 \end{bmatrix} \text{ has } r_{11} = \|\mathbf{a}_1\| = \mathbf{10} \text{ and unit vector } \mathbf{q}_1 = \frac{\mathbf{a}_1}{10}$$

$$r_{12} = \mathbf{q}_1^T \mathbf{a}_2 = \begin{bmatrix} 0.6 & 0.8 \end{bmatrix} \begin{bmatrix} 2 \\ 6 \end{bmatrix} = \mathbf{6} \quad \text{and} \quad A = \begin{bmatrix} 0.6 \\ 0.8 \end{bmatrix} \begin{bmatrix} 10 & 6 \end{bmatrix} + \begin{bmatrix} 0 & -1.6 \\ 0 & 1.2 \end{bmatrix}$$

$$\text{That last column has length } r_{22} = \mathbf{2} \text{ and } A = \begin{bmatrix} 6 & 2 \\ 8 & 6 \end{bmatrix} = \begin{bmatrix} 0.6 & -0.8 \\ 0.8 & 0.6 \end{bmatrix} \begin{bmatrix} \mathbf{10} & \mathbf{6} \\ \mathbf{0} & \mathbf{2} \end{bmatrix}.$$

This product $A = QR$ or $Q^T A = R$ confirms that every $r_{ij} = \mathbf{q}_i^T \mathbf{a}_j$ (row times column). The mysterious matrix R just contains inner products of \mathbf{q} 's and \mathbf{a} 's. R is triangular because \mathbf{q}_i does not involve \mathbf{a}_j for $j > i$. Gram-Schmidt uses only $\mathbf{a}_1, \dots, \mathbf{a}_i$ to construct \mathbf{q}_i .

The next vector \mathbf{q}_2 is the first column of A_2 divided by its length. The next vector \mathbf{r}_2^T contains (after a first zero) the inner products of \mathbf{q}_2 with columns of A_2 :

$$A = \mathbf{q}_1 \mathbf{r}_1^T + \mathbf{q}_2 \mathbf{r}_2^T + \begin{bmatrix} \mathbf{0} & \mathbf{0} & A_3 \end{bmatrix}. \quad (6)$$

All columns of A_3 are orthogonal to \mathbf{q}_1 and \mathbf{q}_2 .

After n steps this is $A = QR$. Only the order of the orthogonalization steps has been modified—by subtracting components (projections) from the columns of A as soon as each new \mathbf{q}_k direction has been found.

Now come the last two factorizations of A .

Eigenvalue Decomposition $A = X \Lambda X^{-1} = X \Lambda Y^T$

The effect of n independent eigenvectors $\mathbf{x}_1, \dots, \mathbf{x}_n$ is to diagonalize the matrix A . Those “right eigenvectors” are the columns of X . Column by column, we see $AX = X\Lambda$. Then $\Lambda = X^{-1}AX$ is the diagonal matrix of eigenvalues, as usual.

To keep the balance between columns and rows, recognize that the rows of X^{-1} are the “left eigenvectors” of A . This is expressed by $X^{-1}A = \Lambda X^{-1}$. Writing $\mathbf{y}_1^T, \dots, \mathbf{y}_n^T$ for the rows of X^{-1} we have $\mathbf{y}_j^T A = \lambda_j \mathbf{y}_j^T$. So the diagonalization $A = X \Lambda X^{-1}$ actually has the more symmetric form $A = X \Lambda Y^T$:

$$\text{Right and left eigenvectors} \quad A = X \Lambda Y^T = \lambda_1 \mathbf{x}_1 \mathbf{y}_1^T + \dots + \lambda_n \mathbf{x}_n \mathbf{y}_n^T \quad (7)$$

Notice that these left eigenvectors \mathbf{y}_i^T are normalized by $Y^T X = X^{-1} X = I$. This requires $\mathbf{y}_j^T \mathbf{x}_j = 1$ and confirms the biorthogonality $\mathbf{y}_i^T \mathbf{x}_j = \delta_{ij}$ of the two sets of eigenvectors. A symmetric matrix has $\mathbf{y}_j = \mathbf{x}_j = \mathbf{q}_j$ and orthonormal eigenvectors. Then $S = Q \Lambda Q^T$.

Singular Value Decomposition $A = U\Sigma V^T$ By comparing with the diagonalization $A = X\Lambda X^{-1} = X\Lambda Y^T$, we see the parallels between a right-left eigenvector decomposition (for a diagonalizable matrix) and a right-left singular value decomposition $A = U\Sigma V^T$ (for any matrix):

$$\text{The SVD with singular vectors} \quad A = U\Sigma V^T = \sigma_1 \mathbf{u}_1 \mathbf{v}_1^T + \cdots + \sigma_r \mathbf{u}_r \mathbf{v}_r^T. \quad (8)$$

For every matrix A , the right singular vectors in V are orthonormal and the left singular vectors $\mathbf{u}_j = A\mathbf{v}_j / \|A\mathbf{v}_j\|$ are orthonormal. Those \mathbf{v} 's and \mathbf{u} 's are eigenvectors of $A^T A$ and AA^T .

$$A^T A \mathbf{v}_j = \sigma_j^2 \mathbf{v}_j \quad \text{and} \quad (AA^T) A \mathbf{v}_j = \sigma_j^2 A \mathbf{v}_j \quad \text{and} \quad (AA^T) \mathbf{u}_j = \sigma_j^2 \mathbf{u}_j. \quad (9)$$

These matrices have the same nonzero eigenvalues $\sigma_1^2, \dots, \sigma_r^2$. The ranks of A and $A^T A$ and AA^T are r . When the singular values are in decreasing order $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_r > 0$, the most important piece of A is $\sigma_1 \mathbf{u}_1 \mathbf{v}_1^T$:

$$\|A\| = \sigma_1 \quad \text{and} \quad \|A - \sigma_1 \mathbf{u}_1 \mathbf{v}_1^T\| = \sigma_2. \quad (10)$$

The rank one matrix closest to A is $\sigma_1 \mathbf{u}_1 \mathbf{v}_1^T$. The difference $A - \sigma_1 \mathbf{u}_1 \mathbf{v}_1^T$ will have singular values $\sigma_2 \geq \sigma_3 \geq \dots \geq \sigma_r$. At every step—not only this first step—the SVD produces the matrix $A_k = \sigma_1 \mathbf{u}_1 \mathbf{v}_1^T + \cdots + \sigma_k \mathbf{u}_k \mathbf{v}_k^T$ of rank k that is **closest to the original A** :

$$\text{(Eckart-Young)} \quad \|A - B\| \geq \|A - A_k\| = \sigma_{k+1} \quad \text{if } B \text{ has rank } k. \quad (11)$$

Thus the SVD produces the rank one pieces $\sigma_i \mathbf{u}_i \mathbf{v}_i^T$ *in order of importance*. This is a central result in data science. We are measuring all these matrices by their spectral norms:

$$\|A\| = \text{maximum of } \|A\mathbf{x}\| = \text{maximum of } \mathbf{u}^T A \mathbf{v} \quad \text{with } \|\mathbf{x}\| = \|\mathbf{u}\| = \|\mathbf{v}\| = 1.$$

In Principal Component Analysis, the leading singular vectors are “principal components”. In statistics, each row of A is centered by subtracting its mean value from its entries. Then $S = AA^T$ is a sample covariance matrix. Its top eigenvector \mathbf{u}_1 represents the combination of rows of S with the greatest variance.

Note. For very large data matrices, the SVD is too expensive to compute. An approximation takes its place. That approximation often uses the inexpensive steps of elimination! Now elimination may begin with the *largest entry* of A , and not necessarily with a_{11} .

Geometrically, the singular values in Σ stretch the unit circle $\|\mathbf{x}\| = 1$ into an ellipse. The factorization $U\Sigma V^T = (\textit{orthogonal})$ times $(\textit{diagonal})$ times $(\textit{orthogonal})$ expresses any matrix (roughly speaking) as a rotation times a stretching times a rotation. This has become central to numerical linear algebra.

It may surprise the reader (as it did the author) that the columns of X in $A = X\Lambda Y^T$ are *right* eigenvectors, while the columns of U in $A = U\Sigma V^T$ are called *left* singular vectors. Perhaps this just confirms that mathematics is a human and fallible (and wonderful) joint enterprise of us all.




Figure 1: U and V are rotations and possible reflections. Σ stretches circle to ellipse.

Factorizations can fail! Of the five principal factorizations, only two are guaranteed. Every symmetric matrix has the form $S = Q\Lambda Q^T$ and every matrix has the form $A = U\Sigma V^T$. The cases of failure are important too (or adjustment more than failure). $A = LU$ now requires an “echelon form E ” and diagonalization needs a “Jordan form J ”. Matrix multiplication is still columns times rows.

Elimination to row reduced echelon form $A = CE = (m \text{ by } r)(r \text{ by } n)$.

The rank of all three matrices is r . E normally comes from operations on the rows of A . Then it may have zero rows. If we work instead with columns of A , the factors C and E have direct meaning.

C contains r independent columns of A (a basis for the column space of A).

E expresses each column of A as a combination of the basic columns in C .

To choose those independent columns, work from left to right ($j = 1$ to $j = n$).

A column of A is *included* in C when it is not a combination of preceding columns.

The r corresponding columns of E contain the r by r identity matrix.

A column of A is *excluded* from C when it is a combination of preceding columns of C .

The corresponding column of E contains the coefficients in that combination.

Example
$$A = \begin{bmatrix} 1 & 4 & 7 \\ 2 & 5 & 8 \\ 3 & 6 & 9 \end{bmatrix} = \begin{bmatrix} 1 & 4 \\ 2 & 5 \\ 3 & 6 \end{bmatrix} \begin{bmatrix} 1 & 0 & -1 \\ 0 & 1 & 2 \end{bmatrix} = CE$$

The entries of E are uniquely determined because C has independent columns.

The remarkable point is that **this E coincides with the row reduced echelon form of A** —except that zero rows are here discarded. Each “1” from the identity matrix inside E is the first nonzero in that row of E . Those 1’s appear in descending order in I and E . The key to the rest of E is this :

The nullspace of E is the nullspace of $A = CE$. (C has independent columns.)

Therefore the row space of E is the row space of A .

Then our E must be the row reduced echelon form without its zero rows.

The nullspace and row space of any matrix are orthogonal complements. So the column construction and the row construction must find the same E .

Gram-Schmidt echelon form $A = QU = (m \text{ by } r)(r \text{ by } n)$.

When the columns of A are linearly dependent ($r < n$), Gram-Schmidt breaks down. Only r orthonormal columns q_j are combinations of columns of A . Those q_j are combinations of the independent columns in C above, because Gram-Schmidt also works left to right. Then $A = CE$ is the same as Gram-Schmidt $C = QR$ multiplied by E :

$$A = CE = (QR)E = Q(RE) = QU. \quad (12)$$

Q has r orthonormal columns: $Q^T Q = I$. The upper triangular $U = RE$ combines an r by r invertible triangular matrix U from Gram-Schmidt and the r by n echelon matrix E from elimination. The nullspaces of U and E and A are all the same, and $A = QU$.

The Jordan form $A = GJG^{-1} = GJH^T$ of a square matrix

Now it is not the columns of A but its eigenvectors that fail to span \mathbf{R}^n . We need to supplement the eigenvectors by “generalized eigenvectors”:

$$Ag_j = \lambda_j g_j \text{ is supplemented as needed by } Ag_k = \lambda_k g_k + g_{k-1}$$

The former puts λ_j on the diagonal of J . The latter produces also a “1” on the superdiagonal. The construction of the Jordan form J is an elegant mess (and I believe that a beginning linear algebra class has more important things to do : the five factorizations).

The only novelty is to see *left generalized eigenvectors* h_i^T when we invert G . Start from a 2 by 2 Jordan block :

$$AG = A \begin{bmatrix} g_1 & g_2 \end{bmatrix} = \begin{bmatrix} \lambda g_1 & \lambda g_2 + g_1 \end{bmatrix} = \begin{bmatrix} g_1 & g_2 \end{bmatrix} \begin{bmatrix} \lambda & 1 \\ 0 & \lambda \end{bmatrix} = GJ. \quad (13)$$

Then $AG = GJ$ gives $G^{-1}A = JG^{-1}$. Write h^T for the rows of G^{-1} (the left generalized eigenvectors):

$$\begin{bmatrix} h_1^T \\ h_2^T \end{bmatrix} A = \begin{bmatrix} \lambda & 1 \\ 0 & \lambda \end{bmatrix} \begin{bmatrix} h_1^T \\ h_2^T \end{bmatrix} \text{ is } h_1^T A = \lambda h_1^T + h_2^T \text{ and } h_2^T A = \lambda h_2^T. \quad (14)$$

In the same way that $A = X\Lambda X^{-1}$ became $A = X\Lambda Y^T$, the Jordan decomposition $A = GJG^{-1}$ has become $A = GJH^T$. We have rows times columns :

$$A = \lambda g_1 h_1^T + (\lambda g_2 + g_1) h_2^T = g_1 (\lambda h_1^T + h_2^T) + \lambda g_2 h_2^T. \quad (15)$$

Summary The first lines of this paper connect inner and outer products. Those are “rows times columns” and “columns times rows”: Level 1 multiplication and Level 3 multiplication.

$$\text{Level 1 Inner product } \mathbf{a}^T \mathbf{b}: \text{ row times column} = \begin{bmatrix} a_1 & \dots & a_n \end{bmatrix} \begin{bmatrix} b_1 \\ \vdots \\ b_n \end{bmatrix} = \text{scalar}$$

$$\text{Level 2 Linear combination } \mathbf{A} \mathbf{b}: \begin{bmatrix} a_1 & \dots & a_n \end{bmatrix} \begin{bmatrix} b_1 \\ \vdots \\ b_n \end{bmatrix} = \sum a_j b_j = \text{vector}$$

$$\text{Level 3 Outer product } \mathbf{a} \mathbf{b}^T: \text{ column times row} \begin{bmatrix} a_1 \\ \vdots \\ a_n \end{bmatrix} \begin{bmatrix} b_1 & \dots & b_n \end{bmatrix} = \text{matrix}$$

The product AB of n by n matrices can be computed at every level, always with the same n^3 multiplications:

$$\text{Level 1 } (n^2 \text{ inner products})(n \text{ multiplications each}) = n^3$$

$$\text{Level 2 } (n \text{ columns } \mathbf{A} \mathbf{b})(n^2 \text{ multiplications each}) = n^3$$

$$\text{Level 3 } (n \text{ outer products})(n^2 \text{ multiplications each}) = n^3$$

These correspond to the three levels of Basic Linear Algebra Subroutines (BLAS). Those are the core operations in LAPACK at the center of computational linear algebra [1]. Our factorizations are produced by four of the most frequently used MATLAB commands: **lu**, **gr**, **eig**, and **svd**

Finally we verify the most important property of matrix multiplication.

The associative law is $(\mathbf{A} \mathbf{B}) \mathbf{C} = \mathbf{A} (\mathbf{B} \mathbf{C})$

Multiplying columns times rows satisfies this fundamental law. The matrices A, B, C are m by n , n by p , and p by q . When $n = p = 1$ and B is a scalar b_{jk} , the laws of arithmetic give two equal matrices of rank one:

$$(\mathbf{a}_j b_{jk}) \mathbf{c}_k^T = \mathbf{a}_j (b_{jk} \mathbf{c}_k^T). \quad (16)$$

The full law $(\mathbf{A} \mathbf{B}) \mathbf{C} = \mathbf{A} (\mathbf{B} \mathbf{C})$ will follow from the agreement of double sums:

$$\sum_{k=1}^p \sum_{j=1}^n (\mathbf{a}_j b_{jk}) \mathbf{c}_k^T = \sum_{j=1}^n \sum_{k=1}^p \mathbf{a}_j (b_{jk} \mathbf{c}_k^T). \quad (17)$$

We are just adding the same np terms. After the inner sum on each side, this becomes

$$\sum_{k=1}^p (\mathbf{A} \mathbf{B})_k \mathbf{c}_k^T = \sum_{j=1}^n \mathbf{a}_j (\mathbf{B} \mathbf{C})_j^T. \quad (18)$$

With another column-row multiplication this is $(\mathbf{A} \mathbf{B}) \mathbf{C} = \mathbf{A} (\mathbf{B} \mathbf{C})$. Parentheses are not needed in $Q \Lambda Q^T$ and $X \Lambda Y^T$ and $U \Sigma V^T$.

REFERENCES

1. E. Anderson et al., *LAPACK Users' Guide*, 3rd edition, SIAM, 1999.
2. G. Strang, *Introduction to Linear Algebra*, 5th edition, Wellesley-Cambridge Press, 2016.
3. L. N. Trefethen and D. Bau, *Numerical Linear Algebra*, SIAM, 1997.