



# THE DECOMPOSITIONAL APPROACH TO MATRIX COMPUTATION

*The introduction of matrix decomposition into numerical linear algebra revolutionized matrix computations. This article outlines the decompositional approach, comments on its history, and surveys the six most widely used decompositions.*

In 1951, Paul S. Dwyer published *Linear Computations*, perhaps the first book devoted entirely to numerical linear algebra.<sup>1</sup> Digital computing was in its infancy, and Dwyer focused on computation with mechanical calculators. Nonetheless, the book was state of the art. Figure 1 reproduces a page of the book dealing with Gaussian elimination. In 1954, Alston S. Householder published *Principles of Numerical Analysis*,<sup>2</sup> one of the first modern treatments of high-speed digital computation. Figure 2 reproduces a page from this book, also dealing with Gaussian elimination.

The contrast between these two pages is striking. The most obvious difference is that Dwyer

used scalar equations whereas Householder used partitioned matrices. But a deeper difference is that while Dwyer started from a system of equations, Householder worked with a (block) LU decomposition—the factorization of a matrix into the product of lower and upper triangular matrices.

Generally speaking, a decomposition is a factorization of a matrix into simpler factors. The underlying principle of the decompositional approach to matrix computation is that it is not the business of the matrix algorithmists to solve particular problems but to construct computational platforms from which a variety of problems can be solved. This approach, which was in full swing by the mid-1960s, has revolutionized matrix computation.

To illustrate the nature of the decompositional approach and its consequences, I begin with a discussion of the Cholesky decomposition and the solution of linear systems—essentially the decomposition that Gauss computed in his elim-

or a non-diagonal pivot, is used. The coefficient serving as a pivot should be different from zero.

By dividing the first equation by  $a_{11}$  and letting  $a_{1i}/a_{11} = b_{1i}$ , as in (6.3.1), the equations (4.1.2) become

$$(1) \quad \begin{aligned} a_{21}x_1 + a_{22}x_2 + a_{23}x_3 + a_{24}x_4 &= a_{25} \\ a_{31}x_1 + a_{32}x_2 + a_{33}x_3 + a_{34}x_4 &= a_{35} \\ a_{41}x_1 + a_{42}x_2 + a_{43}x_3 + a_{44}x_4 &= a_{45} \\ x_1 + b_{12}x_2 + b_{13}x_3 + b_{14}x_4 &= b_{15}. \end{aligned}$$

Multiply the last equation by  $a_{21}$  and subtract from the first equation, by  $a_{31}$  and subtract from the second equation, by  $a_{41}$  and subtract from the third equation, and get the three equations

$$(2) \quad \begin{aligned} g_{22}x_2 + g_{23}x_3 + g_{24}x_4 &= g_{25} \\ g_{32}x_2 + g_{33}x_3 + g_{34}x_4 &= g_{35} \\ g_{42}x_2 + g_{43}x_3 + g_{44}x_4 &= g_{45} \end{aligned}$$

with

$$(3) \quad g_{ij-1} = a_{ij} - a_{i1}b_{1j}.$$

Now

$$(4) \quad \frac{g_{ij-1}}{g_{ii-1}} = \frac{a_{ij} - a_{i1}b_{1j}}{a_{ii} - a_{i1}b_{1i}} = \frac{\frac{a_{ij}}{a_{ii}} - \frac{a_{i1}}{a_{ii}}b_{1j}}{\frac{a_{ii}}{a_{ii}} - \frac{a_{i1}}{a_{ii}}b_{1i}} = \frac{b_{ij} - \frac{a_{i1}}{a_{ii}}b_{1j}}{1 - \frac{a_{i1}}{a_{ii}}b_{1i}} = b_{ij-1}$$

as defined by (6.3.4). We divide the first equation of (2) by  $g_{22}$  and place the results in the bottom row to get

$$(5) \quad \begin{aligned} g_{32}x_2 + g_{33}x_3 + g_{34}x_4 &= g_{35} \\ g_{42}x_2 + g_{43}x_3 + g_{44}x_4 &= g_{45} \\ x_2 + b_{23}x_3 + b_{24}x_4 &= b_{25} \end{aligned}$$

We eliminate as before and obtain

$$(6) \quad \begin{aligned} g_{33}x_3 + g_{34}x_4 &= g_{35} \\ g_{43}x_3 + g_{44}x_4 &= g_{45} \end{aligned}$$

Figure 1. This page from *Linear Computations* shows that Paul Dwyer's approach begins with a system of scalar equations. Courtesy of John Wiley & Sons.

ination algorithm. This article also provides a tour of the five other major matrix decompositions, including the pivoted LU decomposition, the QR decomposition, the spectral decomposition, the Schur decomposition, and the singular value decomposition.

A disclaimer is in order. This article deals primarily with dense matrix computations. Although the decompositional approach has greatly influenced iterative and direct methods for sparse matrices, the ways in which it has affected them are different from what I describe here.

### The Cholesky decomposition and linear systems

We can use the decompositional approach to solve the system

$$(1) \quad \mathbf{Ax} = \mathbf{b}$$

where  $\mathbf{A}$  is positive definite. It is well known that

unknowns is equivalent to the operation of multiplying the system by a particular unit lower triangular matrix—a matrix, in fact, whose off-diagonal non-null elements are all in the same column. The product of all these unit lower triangular matrices is again a unit lower triangular matrix, and hence the entire process of elimination (as opposed to that of back substitution) is equivalent to that of multiplying the system by a suitably chosen unit lower triangular matrix. Since the matrix of the resulting system is clearly upper triangular, these considerations constitute another proof of the possibility of factorizing  $\mathbf{A}$  into a unit lower triangular matrix and an upper triangular matrix.

For the system

$$\mathbf{Ax} = \mathbf{y},$$

after eliminating any one of the variables, the effect to that point is that of having selected a unit lower triangular matrix of the form

$$\begin{pmatrix} L_{11} & 0 \\ & I_{22} \end{pmatrix}$$

where  $L_{11}$  is itself unit lower triangular, in such a way that  $\mathbf{A}$  is factored

$$(2.21.1) \quad \mathbf{A} = \begin{pmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{pmatrix} = \begin{pmatrix} L_{11} & 0 \\ & I_{22} \end{pmatrix} \begin{pmatrix} W_{11} & W_{12} \\ 0 & M_{22} \end{pmatrix},$$

with  $W_{11}$  upper triangular but  $M_{22}$  not. Hence

$$(2.21.2) \quad M_{22} = A_{22} - A_{21}A_{11}^{-1}A_{12}.$$

The original system has at this stage been replaced by the system

$$(2.21.3) \quad \begin{pmatrix} W_{11} & W_{12} \\ 0 & M_{22} \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = \begin{pmatrix} y_1 \\ y_2 \end{pmatrix}$$

where

$$(2.21.4) \quad \begin{pmatrix} L_{11} & 0 \\ & I_{22} \end{pmatrix} \begin{pmatrix} z_1 \\ z_2 \end{pmatrix} = \begin{pmatrix} y_1 \\ y_2 \end{pmatrix} = \mathbf{y}.$$

The matrices  $L_{11}$  and  $L_{21}$  are not themselves written down. The partial system

$$M_{22}z_2 = z_2$$

represents those equations from which further elimination remains to be done, and this can be treated independently of the other equations of the system, which fact explains why it is unnecessary to obtain the  $L$  matrices explicitly.

If the upper left-hand element of  $M_{22}$  vanishes, this cannot be used in the next step of the elimination, and it is not advantageous to use it when it is small. Hence rows or columns, or both, in  $M_{22}$  must be

Figure 2. On this page from *Principles of Numerical Analysis*, Alston Householder uses partitioned matrices and LU decomposition. Courtesy of McGraw-Hill.

$\mathbf{A}$  can be factored in the form

$$(2) \quad \mathbf{A} = \mathbf{R}^T \mathbf{R}$$

where  $\mathbf{R}$  is an upper triangular matrix. The factorization is called the *Cholesky decomposition* of  $\mathbf{A}$ .

The factorization in Equation 2 can be used to solve linear systems as follows. If we write the system in the form  $\mathbf{R}^T \mathbf{R}x = \mathbf{b}$  and set  $\mathbf{y} = \mathbf{R}^T \mathbf{b}$ , then  $x$  is the solution of the triangular system  $\mathbf{R}x = \mathbf{y}$ . However, by definition  $\mathbf{y}$  is the solution of the system  $\mathbf{R}^T \mathbf{y} = \mathbf{b}$ . Consequently, we have reduced the problem to the solution of two triangular systems, as illustrated in the following algorithm:

1. Solve the system  $\mathbf{R}^T \mathbf{y} = \mathbf{b}$ .
2. Solve the system  $\mathbf{R}x = \mathbf{y}$ .

Because triangular systems are easy to solve, the introduction of the Cholesky decomposition has




Figure 3. These varieties of Gaussian elimination are all numerically equivalent.

transformed our problem into one for which the solution can be readily computed.

We can use such decompositions to solve more than one problem. For example, the following algorithm solves the system  $\mathbf{A}^T \mathbf{x} = \mathbf{b}$ :

1. Solve the system  $\mathbf{R}\mathbf{y} = \mathbf{b}$ .
2. Solve the system  $\mathbf{R}^T \mathbf{x} = \mathbf{y}$ . (4)

Again, in many statistical applications we want to compute the quantity  $\rho = \mathbf{x}^T \mathbf{A}^{-1} \mathbf{x}$ . Because

$$\mathbf{x}^T \mathbf{A}^{-1} \mathbf{x} = \mathbf{x}^T (\mathbf{R}^T \mathbf{R})^{-1} \mathbf{x} = (\mathbf{R}^{-T} \mathbf{x})^T (\mathbf{R}^{-T} \mathbf{x}) \quad (5)$$

we can compute  $\rho$  as follows:

1. Solve the system  $\mathbf{R}^T \mathbf{y} = \mathbf{x}$ .
2.  $\rho = \mathbf{y}^T \mathbf{y}$ . (6)

The decompositional approach can also save computation. For example, the Cholesky decomposition requires  $O(n^3)$  operations to compute, whereas the solution of triangular systems requires only  $O(n^2)$  operations. Thus, if you recognize that a Cholesky decomposition is being

used to solve a system at one point in a computation, you can reuse the decomposition to do the same thing later without having to recompute it. Historically, Gaussian elimination and its variants (including Cholesky's algorithm) have solved the system in Equation 1 by reducing it to an equivalent triangular system. This mixes the computation of the decomposition with the solution of the first triangular system in Equation 3, and it is not obvious how to reuse the elimination when a new right-hand side presents itself. A naive programmer is in danger of performing the reduction from the beginning, thus repeating the lion's share of the work. On the other hand, a program that knows a decomposition is in the background can reuse it as needed.

(By the way, the problem of recomputing decompositions has not gone away. Some matrix packages hide the fact that they repeatedly compute a decomposition by providing drivers to solve linear systems with a call to a single routine. If the program calls the routine again with the same matrix, it recomputes the decomposition—unnecessarily. Interpretive matrix systems such as Matlab and Mathematica have the same problem—they hide decompositions behind operators and function calls. Such are the consequences of not stressing the decompositional approach to the consumers of matrix algorithms.)

Another advantage of working with decompositions is unity. There are different ways of organizing the operations involved in solving linear systems by Gaussian elimination in general and Cholesky's algorithm in particular. Figure 3 illustrates some of these arrangements: a white area contains elements from the original matrix, a dark area contains the factors, a light gray area contains partially processed elements, and the boundary strips contain elements about to be processed. Most of these variants were originally presented in scalar form as new algorithms. Once you recognize that a decomposition is involved, it is easy to see the essential unity of the various algorithms.

All the variants in Figure 3 are numerically equivalent. This means that one rounding-error analysis serves all. For example, the Cholesky algorithm, in whatever guise, is backward stable: the computed factor  $\mathbf{R}$  satisfies

$$(\mathbf{A} + E) = \mathbf{R}^T \mathbf{R} \quad (7)$$

where  $E$  is of the size of the rounding unit relative to  $\mathbf{A}$ . Establishing this backward is usually the most difficult part of an analysis of the use

of a decomposition to solve a problem. For example, once Equation 7 has been established, the rounding errors involved in the solutions of the triangular systems in Equation 3 can be incorporated in  $E$  with relative ease. Thus, another advantage of the decompositional approach is that it concentrates the most difficult aspects of rounding-error analysis in one place.

In general, if you change the elements of a positive definite matrix, you must recompute its Cholesky decomposition from scratch. However, if the change is structured, it may be possible to compute the new decomposition directly from the old—a process known as *updating*. For example, you can compute the Cholesky decomposition of  $A + xx^T$  from that of  $A$  in  $O(n^2)$  operations, an enormous savings over the *ab initio* computation of the decomposition.

Finally, the decompositional approach has greatly affected the development of software for matrix computation. Instead of wasting energy developing code for a variety of specific applications, the producers of a matrix package can concentrate on the decompositions themselves, perhaps with a few auxiliary routines to handle the most important applications. This approach has informed the major public-domain packages: the Handbook series,<sup>3</sup> Eispack,<sup>4</sup> Linpack,<sup>5</sup> and Lapack.<sup>6</sup> A consequence of this emphasis on decompositions is that software developers have found that most algorithms have broad computational features in common—features that can be relegated to basic linear-algebra subprograms (such as Blas), which can then be optimized for specific machines.<sup>7-9</sup>

For easy reference, the sidebar “Benefits of the decompositional approach” summarizes the advantages of decomposition.

## History

All the widely used decompositions had made their appearance by 1909, when Schur introduced the decomposition that now bears his name. However, with the exception of the Schur decomposition, they were not cast in the language of matrices (in spite of the fact that matrices had been introduced in 1858<sup>10</sup>). I provide some historical background for the individual decompositions later, but it is instructive here to consider how the originators proceeded in the absence of matrices.

Gauss, who worked with positive definite systems defined by the normal equations for least squares, described his elimination procedure as

## Benefits of the decompositional approach

- A matrix decomposition solves not one but many problems.
- A matrix decomposition, which is generally expensive to compute, can be reused to solve new problems involving the original matrix.
- The decompositional approach often shows that apparently different algorithms are actually computing the same object.
- The decompositional approach facilitates rounding-error analysis.
- Many matrix decompositions can be updated, sometimes with great savings in computation.
- By focusing on a few decompositions instead of a host of specific problems, software developers have been able to produce highly effective matrix packages.

the reduction of a quadratic form  $\varphi(x) = x^T A x$  (I am simplifying a little here). In terms of the Cholesky factorization  $A = R^T R$ , Gauss wrote  $\varphi(x)$  in the form

$$\begin{aligned} \varphi(x) &= (r_1^T x)^2 + (r_2^T x)^2 + \dots + (r_n^T x)^2 \\ &\equiv \rho_1^2(x) + \rho_2^2(x) + \dots + \rho_n^2(x) \end{aligned} \quad (8)$$

where  $r_i^T$  is the  $i$ th row of  $R$ . Thus Gauss reduced  $\varphi(x)$  to a sum of squares of linear functions  $\rho_i$ . Because  $R$  is upper triangular, the function  $\rho_i(x)$  depends only on the components  $x_i, \dots, x_n$  of  $x$ . Since the coefficients in the linear forms  $\rho_i$  are the elements of  $R$ , Gauss, by showing how to compute the  $\rho_i$ , effectively computed the Cholesky decomposition of  $A$ .

Other decompositions were introduced in other ways. For example, Jacobi introduced the LU decomposition as a decomposition of a bilinear form into a sum of products of linear functions having an appropriate triangularity with respect to the variables. The singular value decomposition made its appearance as an orthogonal change of variables that diagonalized a bilinear form. Eventually, all these decompositions found expressions as factorizations of matrices.<sup>11</sup>

The process by which decomposition became so important to matrix computations was slow and incremental. Gauss certainly had the spirit. He used his decomposition to perform many tasks, such as computing variances, and even used it to update least-squares solutions. But Gauss never regarded his decomposition as a matrix factorization, and it would be anachro-

nistic to consider him the father of the decompositional approach.

In the 1940s, awareness grew that the usual algorithms for solving linear systems involved matrix factorization.<sup>12,13</sup> John Von Neumann and H.H. Goldstine, in their ground-breaking error analysis of the solution of linear systems, pointed out the division of labor between computing a factorization and solving the system:<sup>14</sup>

We may therefore interpret the elimination method as one which bases the inverting of an arbitrary matrix  $\mathbf{A}$  on the combination of two tricks: First it decomposes  $\mathbf{A}$  into the product of two semi-diagonal matrices  $\mathbf{C}$ ,  $\mathbf{B}$  ..., and consequently the inverse of  $\mathbf{A}$  obtains immediately from those of  $\mathbf{C}$  and  $\mathbf{B}$ . Second it forms their inverses by a simple, explicit, inductive process.

In the 1950s and early 1960s, Householder systematically explored the relation between various algorithms in matrix terms. His book *The Theory of Matrices in Numerical Analysis* is the mathematical epitome of the decompositional approach.<sup>15</sup>

In 1954, Givens showed how to reduce a symmetric matrix  $\mathbf{A}$  to tridiagonal form by orthogonal transformation.<sup>16</sup> The reduction was merely a way station to the computation of the eigenvalues of  $\mathbf{A}$ , and at the time no one thought of it as a decomposition. However, it and other intermediate forms have proven useful in their own right and have become a staple of the decompositional approach.

In 1961, James Wilkinson gave the first backward rounding-error analysis of the solutions of linear systems.<sup>17</sup> Here, the division of labor is complete. He gives one analysis of the computation of the LU decomposition and another of the solution of triangular systems and then combines the two. Wilkinson continued analyzing various algorithms for computing decompositions, introducing uniform techniques for dealing with the transformations used in the computations. By the time his book *Algebraic Eigenvalue Problem*<sup>18</sup> appeared in 1965, the decompositional approach was firmly established.

## The big six

There are many matrix decompositions, old and new, and the list of the latter seems to grow daily. Nonetheless, six decompositions hold the center. The reason is that they are useful and stable—they have important applications and the

algorithms that compute them have a satisfactory backward rounding-error analysis (see Equation 7). In this brief tour, I provide references only for details that cannot be found in the many excellent texts and monographs on numerical linear algebra,<sup>18–26</sup> the Handbook series,<sup>3</sup> or the *LINPACK Users' Guide*.<sup>5</sup>

## The Cholesky decomposition

**Description.** Given a positive definite matrix  $\mathbf{A}$ , there is a unique upper triangular matrix  $\mathbf{R}$  with positive diagonal elements such that

$$\mathbf{A} = \mathbf{R}^T \mathbf{R}.$$

In this form, the decomposition is known as the Cholesky decomposition. It is often written in the form

$$\mathbf{A} = \mathbf{L} \mathbf{D} \mathbf{L}^T$$

where  $\mathbf{D}$  is diagonal and  $\mathbf{L}$  is unit lower triangular (that is,  $\mathbf{L}$  is lower triangular with ones on the diagonal).

**Applications.** The Cholesky decomposition is used primarily to solve positive definite linear systems, as in Equations 3 and 6. It can also be employed to compute quantities useful in statistics, as in Equation 4.

**Algorithms.** A Cholesky decomposition can be computed using any of the variants of Gaussian elimination (see Figure 3)—modified, of course, to take advantage of symmetry. All these algorithms take approximately  $n^3/6$  floating-point additions and multiplications. The algorithm Cholesky proposed corresponds to the diagram in the lower right of Figure 3.

**Updating.** Given a Cholesky decomposition  $\mathbf{A} = \mathbf{R}^T \mathbf{R}$ , you can calculate the Cholesky decomposition of  $\mathbf{A} + \mathbf{x}\mathbf{x}^T$  from  $\mathbf{R}$  and  $\mathbf{x}$  in  $O(n^2)$  floating-point additions and multiplications. The Cholesky decomposition of  $\mathbf{A} - \mathbf{x}\mathbf{x}^T$  can be calculated with the same number of operations. The latter process, which is called *downdating*, is numerically less stable than updating.

**The pivoted Cholesky decomposition.** If  $\mathbf{P}$  is a permutation matrix and  $\mathbf{A}$  is positive definite, then  $\mathbf{P}^T \mathbf{A} \mathbf{P}$  is said to be a diagonal permutation of  $\mathbf{A}$  (among other things, it permutes the diagonals of  $\mathbf{A}$ ). Any diagonal permutation of  $\mathbf{A}$  is positive definite and has a Cholesky factor. Such a factorization is called a *pivoted Cholesky factorization*. There are many ways to pivot a Cholesky decomposition, but the most common one produces a factor satisfying

$$r_{kk}^2 \geq \sum_{j>k} \sum_{i \geq k} r_{ij}^2 \quad (9)$$

In particular, if  $\mathbf{A}$  is positive semidefinite, this strategy will assure that  $\mathbf{R}$  has the form

$$\mathbf{R} = \begin{pmatrix} R_{11} & R_{12} \\ 0 & 0 \end{pmatrix}$$

where  $R_{11}$  is nonsingular and has the same order as the rank of  $\mathbf{A}$ . Hence, the pivoted Cholesky decomposition is widely used for rank determination.

**History.** The Cholesky decomposition (more precisely, an  $\mathbf{LDL}^T$  decomposition) was the decomposition of Gauss's elimination algorithm, which he sketched in 1809<sup>27</sup> and presented in full in 1810.<sup>28</sup> Benoît published Cholesky's variant posthumously in 1924.<sup>29</sup>

#### The pivoted LU decomposition

**Description.** Given a matrix  $\mathbf{A}$  of order  $n$ , there are permutations  $\mathbf{P}$  and  $\mathbf{Q}$  such that

$$\mathbf{P}^T \mathbf{A} \mathbf{Q} = \mathbf{L} \mathbf{U}$$

where  $\mathbf{L}$  is unit lower triangular and  $\mathbf{U}$  is upper triangular. The matrices  $\mathbf{P}$  and  $\mathbf{Q}$  are not unique, and the process of selecting them is known as *pivoting*.

**Applications.** Like the Cholesky decomposition, the LU decomposition is used primarily for solving linear systems. However, since  $\mathbf{A}$  is a general matrix, this application covers a wide range. For example, the LU decomposition is used to compute the steady-state vector of Markov chains and, with the inverse power method, to compute eigenvectors.

**Algorithms.** The basic algorithm for computing LU decompositions is a generalization of Gaussian elimination to nonsymmetric matrices. When and how this generalization arose is obscure (see Dwyer<sup>1</sup> for comments and references). Except for special matrices (such as positive definite and diagonally dominant matrices), the method requires some form of pivoting for stability. The most common form is partial pivoting, in which pivot elements are chosen from the column to be eliminated. This algorithm requires about  $n^3/3$  additions and multiplications.

Certain contrived examples show that Gaussian elimination with partial pivoting can be unstable. Nonetheless, it works well for the overwhelming majority of real-life problems.<sup>30,31</sup> Why is an open question.

**History.** In establishing the existence of the LU decomposition, Jacobi showed<sup>32</sup> that under certain conditions a bilinear form  $\phi(x, y)$  can be written in the form

$$\phi(x, y) = \rho_1(x)\sigma_1(y) + \rho_2(x)\sigma_2(y) + \dots + \rho_n(x)\sigma_n(y)$$

where  $\rho_i$  and  $\sigma_i$  are linear functions that depend only on the last  $(n - i + 1)$  components of their arguments. The coefficients of the functions are the elements of  $\mathbf{L}$  and  $\mathbf{U}$ .

#### The QR decomposition

**Description.** Let  $\mathbf{A}$  be an  $m \times n$  matrix with  $m \geq n$ . There is an orthogonal matrix  $\mathbf{Q}$  such that

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

where  $\mathbf{R}$  is upper triangular with nonnegative diagonal elements (or positive diagonal elements if  $\mathbf{A}$  is of rank  $n$ ).

If we partition  $\mathbf{Q}$  in the form

$$\mathbf{Q} = (\mathbf{Q}_A \mathbf{Q}_\perp)$$

where  $\mathbf{Q}_A$  has  $n$  columns, then we can write

$$\mathbf{A} = \mathbf{Q}_A \mathbf{R}. \quad (10)$$

This is sometimes called the *QR factorization* of  $\mathbf{A}$ .

**Applications.** When  $\mathbf{A}$  is of rank  $n$ , the columns of  $\mathbf{Q}_A$  form an orthonormal basis for the column space  $\mathcal{R}(\mathbf{A})$  of  $\mathbf{A}$ , and the columns of  $\mathbf{Q}_\perp$  form an orthonormal basis of the orthogonal complement of  $\mathcal{R}(\mathbf{A})$ . In particular,  $\mathbf{Q}_A \mathbf{Q}_A^T$  is the orthogonal projection onto  $\mathcal{R}(\mathbf{A})$ . For this reason, the QR decomposition is widely used in applications with a geometric flavor, especially least squares.

**Algorithms.** There are two distinct classes of algorithms for computing the QR decomposition: Gram–Schmidt algorithms and orthogonal triangularization.

Gram–Schmidt algorithms proceed stepwise by orthogonalizing the  $k$ th columns of  $\mathbf{A}$  against the first  $(k - 1)$  columns of  $\mathbf{Q}$  to get the  $k$ th column of  $\mathbf{Q}$ . There are two forms of the Gram–Schmidt algorithm, the classical and the modified, and they both compute only the factorization in Equation 10. The classical Gram–Schmidt is unstable. The modified form can produce a matrix  $\mathbf{Q}_A$  whose columns deviate from orthogonality. But the deviation is

bounded, and the computed factorization can be used in certain applications—notably computing least-squares solutions. If the orthogonalization step is repeated at each stage—a process known as *reorthogonalization*—both algorithms will produce a fully orthogonal factorization. When  $n \gg m$ , the algorithms without reorthogonalization require about  $mn^2$  additions and multiplications.

The method of orthogonal triangularization proceeds by premultiplying  $\mathbf{A}$  by certain simple orthogonal matrices until the elements below the diagonal are zero. The product of the orthogonal matrices is  $\mathbf{Q}$ , and  $\mathbf{R}$  is the upper triangular part of the reduced  $\mathbf{A}$ . Again, there are two versions. The first reduces the matrix by Householder transformations. The method has the advantage that it represents the entire matrix  $\mathbf{Q}$  in the same amount of memory that is required to hold  $\mathbf{A}$ , a great savings when  $n \gg p$ . The second method reduces  $\mathbf{A}$  by plane rotations. It is less efficient than the first method, but is better suited for matrices with structured patterns of nonzero elements.

#### Relation to the Cholesky decomposition.

From Equation 10, it follows that

$$\mathbf{A}^T \mathbf{A} = \mathbf{R}^T \mathbf{R}. \quad (11)$$

In other words, the triangular factor of the QR decomposition of  $\mathbf{A}$  is the triangular factor of the Cholesky decomposition of the cross-product matrix  $\mathbf{A}^T \mathbf{A}$ . Consequently, many problems—particularly least-squares problems—can be solved using either a QR decomposition from a least-squares matrix or the Cholesky decomposition from the normal equation. The QR decomposition usually gives more accurate results, whereas the Cholesky decomposition is often faster.

**Updating.** Given a QR factorization of  $\mathbf{A}$ , there are stable, efficient algorithms for recomputing the QR factorization after rows and columns have been added to or removed from  $\mathbf{A}$ . In addition, the QR decomposition of the rank-one modification  $\mathbf{A} + xy^T$  can be stably updated.

**The pivoted QR decomposition.** If  $\mathbf{P}$  is a permutation matrix, then  $\mathbf{AP}$  is a permutation of the columns of  $\mathbf{A}$ , and  $(\mathbf{AP})^T(\mathbf{AP})$  is a diagonal permutation of  $\mathbf{A}^T \mathbf{A}$ . In view of the relation of the QR and the Cholesky decompositions, it is not surprising that there is a pivoted QR factorization whose triangular factor  $\mathbf{R}$  satisfies Equation 9. In particular, if  $\mathbf{A}$  has rank  $k$ , then its piv-

oted QR factorization has the form

$$\mathbf{AP} = (\mathbf{Q}_1 \mathbf{Q}_2) \begin{pmatrix} \mathbf{R}_{11} & \mathbf{R}_{12} \\ \mathbf{0} & \mathbf{0} \end{pmatrix}.$$

It follows that either  $\mathbf{Q}_1$  or the first  $k$  columns of  $\mathbf{AP}$  form a basis for the column space of  $\mathbf{A}$ . Thus, the pivoted QR decomposition can be used to extract a set of linearly independent columns from  $\mathbf{A}$ .

**History.** The QR factorization first appeared in a work by Erhard Schmidt on integral equations.<sup>33</sup> Specifically, Schmidt showed how to orthogonalize a sequence of functions by what is now known as the Gram–Schmidt algorithm. (Curiously, Laplace produced the basic formulas<sup>34</sup> but had no notion of orthogonality.) The name QR comes from the QR algorithm, named by Francis (see the history notes for the Schur algorithm, discussed later). Householder introduced Householder transformations to matrix computations and showed how they could be used to triangularize a general matrix.<sup>35</sup> Plane rotations were introduced by Givens,<sup>16</sup> who used them to reduce a symmetric matrix to tridiagonal form. Bogert and Burris appear to be the first to use them in orthogonal triangularization.<sup>36</sup> The first updating algorithm (adding a row) is due to Golub,<sup>37</sup> who also introduced the idea of pivoting.

#### The spectral decomposition

**Description.** Let  $\mathbf{A}$  be a symmetric matrix of order  $n$ . There is an orthogonal matrix  $\mathbf{V}$  such that

$$\mathbf{A} = \mathbf{V} \mathbf{\Lambda} \mathbf{V}^T, \quad \mathbf{\Lambda} = \text{diag}(\lambda_1, \dots, \lambda_n). \quad (12)$$

If  $v_i$  denotes the  $i$ th column of  $\mathbf{V}$ , then  $\mathbf{A}v_i = \lambda_i v_i$ . Thus  $(\lambda_i, v_i)$  is an eigenpair of  $\mathbf{A}$ , and the *spectral decomposition* shown in Equation 12 exhibits the eigenvalues of  $\mathbf{A}$  along with complete orthonormal system of eigenvectors.

**Applications.** The spectral decomposition finds applications wherever the eigensystem of a symmetric matrix is needed, which is to say in virtually all technical disciplines.

**Algorithms.** There are three classes of algorithms to compute the spectral decomposition: the QR algorithm, the divide-and-conquer algorithm, and Jacobi's algorithm. The first two require a preliminary reduction to tridiagonal form by orthogonal similarities. I discuss the QR algorithm in the next section on the Schur decomposition. The divide-and-conquer algo-

algorithm<sup>38,39</sup> is comparatively recent and is usually faster than the QR algorithm when both eigenvalues and eigenvectors are desired; however, it is not suitable for problems in which the eigenvalues vary widely in magnitude. The Jacobi algorithm is much slower than the other two, but for positive definite matrices it may be more accurate.<sup>40</sup> All these algorithms require  $O(n^3)$  operations.

**Updating.** The spectral decomposition can be updated. Unlike the Cholesky and QR decompositions, the algorithm does not result in a reduction in the order of the work—it still remains  $O(n^3)$ , although the order constant is lower.

**History.** The spectral decomposition dates back to an 1829 paper by Cauchy,<sup>41</sup> who introduced the eigenvectors as solutions of equations of the form  $\mathbf{A}\mathbf{x} = \lambda\mathbf{x}$  and proved the orthogonality of eigenvectors belonging to distinct eigenvalues. In 1846, Jacobi<sup>42</sup> gave his famous algorithm for spectral decomposition, which iteratively reduces the matrix in question to diagonal form by a special type of plane rotations, now called Jacobi rotations. The reduction to tridiagonal form by plane rotations is due to Givens<sup>16</sup> and by Householder transformations to Householder.<sup>43</sup>

### The Schur decomposition

**Description.** Let  $\mathbf{A}$  be a matrix of order  $n$ . There is a unitary matrix  $\mathbf{U}$  such that

$$\mathbf{A} = \mathbf{UTU}^H$$

where  $\mathbf{T}$  is upper triangular and  $H$  means conjugate transpose. The diagonal elements of  $\mathbf{T}$  are the eigenvalues of  $\mathbf{A}$ , which, by appropriate choice of  $\mathbf{U}$ , can be made to appear in any order. This decomposition is called a *Schur decomposition of  $\mathbf{A}$* .

A real matrix can have complex eigenvalues and hence a complex Schur form. By allowing  $\mathbf{T}$  to have real  $2 \times 2$  blocks on its diagonal that contain its complex eigenvalues, the entire decomposition can be made real. This is sometimes called a *real Schur form*.

**Applications.** An important use of the Schur form is as an intermediate form from which the eigenvalues and eigenvectors of a matrix can be computed. On the other hand, the Schur decomposition can often be used in place of a complete system of eigenpairs, which, in fact, may not exist. A good example is the solution of Sylvester's equation and its relatives.<sup>44,45</sup>

**Algorithms.** After a preliminary reduction to

Hessenberg form, which is usually done with Householder transformations, the Schur form is computed using the QR algorithm.<sup>46</sup> Elsewhere in this issue, Beresford Parlett discusses the modern form of the algorithm. It is one of the most flexible algorithms in the repertoire, having variants for the spectral decomposition, the singular values decomposition, and the generalized eigenvalue problem.

**History.** Schur introduced his decomposition in 1909.<sup>47</sup> It was the only one of the big six to have been derived in terms of matrices. It was largely ignored until Francis's QR algorithm pushed it into the limelight.

### The singular value decomposition

**Description.** Let  $\mathbf{A}$  be an  $m \times n$  matrix with  $m \geq n$ . There are orthogonal matrices  $\mathbf{U}$  and  $\mathbf{V}$  such that

$$\mathbf{U}^T \mathbf{A} \mathbf{V} = \begin{pmatrix} \Sigma \\ 0 \end{pmatrix}$$

where

$$\Sigma = \text{diag}(\sigma_1, \dots, \sigma_n), \quad \sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_n \geq 0.$$

This decomposition is called the *singular value decomposition of  $\mathbf{A}$* . If  $\mathbf{U}_A$  consists of the first  $n$  columns of  $\mathbf{U}$ , we can write

$$\mathbf{A} = \mathbf{U}_A \Sigma \mathbf{V}^T \tag{13}$$

which is sometimes called the *singular value factorization of  $\mathbf{A}$* .

The diagonal elements of  $\sigma$  are called the *singular values of  $\mathbf{A}$* . The corresponding columns of  $\mathbf{U}$  and  $\mathbf{V}$  are called *left and right singular vectors of  $\mathbf{A}$* .

**Applications.** Most of the applications of the QR decomposition can also be handled by the singular value decomposition. In addition, the singular value decomposition gives a basis for the row space of  $\mathbf{A}$  and is more reliable in determining rank. It can also be used to compute optimal low-rank approximations and to compute angles between subspaces.

**Relation to the spectral decomposition.** The singular value factorization is related to the spectral decomposition in much the same way as the QR factorization is related to the Cholesky decomposition. Specifically, from Equation 13 it follows that

$$\mathbf{A}^T \mathbf{A} = \mathbf{V} \Sigma^2 \mathbf{V}^T.$$


Thus the eigenvalues of the cross-product ma-



trix  $\mathbf{A}^T\mathbf{A}$  are the squares of the singular vectors of  $\mathbf{A}$  and the eigenvectors of  $\mathbf{A}^T\mathbf{A}$  are right singular vectors of  $\mathbf{A}$ .

**Algorithms.** As with the spectral decomposition, there are three classes of algorithms for computing the singular value decomposition: the QR algorithm, a divide-and-conquer algorithm, and a Jacobi-like algorithm. The first two require a reduction of  $\mathbf{A}$  to bidiagonal form. The divide-and-conquer algorithm<sup>49</sup> is often faster than the QR algorithm, and the Jacobi algorithm is the slowest.

**History.** The singular value decomposition was introduced independently by Beltrami in 1873<sup>50</sup> and Jordan in 1874.<sup>51</sup> The reduction to bidiagonal form is due to Golub and Kahan,<sup>52</sup> as is the variant of the QR algorithm. The first Jacobi-like algorithm for computing the singular value decomposition was given by Kogbetliantz.<sup>53</sup>

The big six are not the only decompositions in use; in fact, there are many more. As mentioned earlier, certain intermediate forms—such as tridiagonal and Hessenberg forms—have come to be regarded as decompositions in their own right. Since the singular value decomposition is expensive to compute and not readily updated, rank-revealing alternatives have received considerable attention.<sup>54,55</sup> There are also generalizations of the singular value decomposition and the Schur decomposition for pairs of matrices.<sup>56,57</sup> All crystal balls become cloudy when they look to the future, but it seems safe to say that as long as new matrix problems arise, new decompositions will be devised to solve them. 

## Acknowledgment

This work was supported by the National Science Foundation under Grant No. 970909-8562.

## References

1. P.S. Dwyer, *Linear Computations*, John Wiley & Sons, New York, 1951.
2. A.S. Householder, *Principles of Numerical Analysis*, McGraw-Hill, New York, 1953.
3. J.H. Wilkinson and C. Reinsch, *Handbook for Automatic Computation, Vol. II, Linear Algebra*, Springer-Verlag, New York, 1971.
4. B.S. Garbow et al., "Matrix Eigensystem Routines—Eispack Guide Extension," *Lecture Notes in Computer Science*, Springer-Verlag, New York, 1977.
5. J.J. Dongarra et al., *LINPACK User's Guide*, SIAM, Philadelphia, 1979.
6. E. Anderson et al., *LAPACK Users' Guide*, second ed., SIAM, Philadelphia, 1995.
7. J.J. Dongarra et al., "An Extended Set of Fortran Basic Linear Algebra Subprograms," *ACM Trans. Mathematical Software*, Vol. 14, 1988, pp. 1–17.
8. J.J. Dongarra et al., "A Set of Level 3 Basic Linear Algebra Subprograms," *ACM Trans. Mathematical Software*, Vol. 16, 1990, pp. 1–17.
9. C.L. Lawson et al., "Basic Linear Algebra Subprograms for Fortran Usage," *ACM Trans. Mathematical Software*, Vol. 5, 1979, pp. 308–323.
10. A. Cayley, "A Memoir on the Theory of Matrices," *Philosophical Trans. Royal Soc. of London*, Vol. 148, 1858, pp. 17–37.
11. C.C. Mac Duffee, *The Theory of Matrices*, Chelsea, New York, 1946.
12. P.S. Dwyer, "A Matrix Presentation of Least Squares and Correlation Theory with Matrix Justification of Improved Methods of Solution," *Annals Math. Statistics*, Vol. 15, 1944, pp. 82–89.
13. H. Jensen, *An Attempt at a Systematic Classification of Some Methods for the Solution of Normal Equations*, Report No. 18, Geodætisk Institut, Copenhagen, 1944.
14. J. von Neumann and H.H. Goldstine, "Numerical Inverting of Matrices of High Order," *Bull. Am. Math. Soc.*, Vol. 53, 1947, pp. 1021–1099.
15. A.S. Householder, *The Theory of Matrices in Numerical Analysis*, Dover, New York, 1964.
16. W. Givens, *Numerical Computation of the Characteristic Values of a Real Matrix*, Tech. Report 1574, Oak Ridge Nat'l Laboratory, Oak Ridge, Tenn., 1954.
17. J.H. Wilkinson, "Error Analysis of Direct Methods of Matrix Inversion," *J. ACM*, Vol. 8, 1961, pp. 281–330.
18. J.H. Wilkinson, *The Algebraic Eigenvalue Problem*, Clarendon Press, Oxford, U.K., 1965.
19. Å. Björck, *Numerical Methods for Least Squares Problems*, SIAM, Philadelphia, 1996.
20. B.N. Datta, *Numerical Linear Algebra and Applications*, Brooks/Cole, Pacific Grove, Calif., 1995.
21. J.W. Demmel, *Applied Numerical Linear Algebra*, SIAM, Philadelphia, 1997.
22. N.J. Higham, *Accuracy and Stability of Numerical Algorithms*, SIAM, Philadelphia, 1996.
23. B.N. Parlett, *The Symmetric Eigenvalue Problem*, Prentice-Hall, Englewood Cliffs, N.J., 1980; reissued with revisions by SIAM, Philadelphia, 1998.
24. G.W. Stewart, *Introduction to Matrix Computations*, Academic Press, New York, 1973.
25. G.W. Stewart, *Matrix Algorithms I: Basic Decompositions*, SIAM, Philadelphia, 1998.
26. D.S. Watkins, *Fundamentals of Matrix Computations*, John Wiley & Sons, New York, 1991.
27. C.F. Gauss, *Theoria Motus Corporum Coelestium in Sectionibus Conicis Solem Ambientium [Theory of the Motion of the Heavenly Bodies Moving about the Sun in Conic Sections]*, Perthes and Besser, Hamburg, Germany, 1809.
28. C.F. Gauss, "Disquisitio de elementis ellipticis Palladis [Disquisition on the Elliptical Elements of Pallas]," *Commentationes societatis regiae scientiarum Göttingensis recentiores*, Vol. 1, 1810.
29. C. Benoît, "Note sur une méthode de résolution des équations normales provenant de l'application de la méthode des moindres carrés a un système d'équations linéaires en nombre inférieur à celui des inconnues. — application de la méthode à la résolution d'un système défini d'équations linéaires (Procédé du Commandant Cholesky) [Note on a Method for Solving the Normal Equations Arising from the Application of the Method of Least Squares to a System of Linear Equations whose Number Is Less than the Number of Unknowns—Application of the Method to the Solution of a Definite System of Linear Equations]," *Bulletin Géodésique (Toulouse)*, Vol. 2, 1924, pp. 5–77.

30. L.N. Trefethen and R.S. Schreiber, "Average-Case Stability of Gaussian Elimination," *SIAM J. Matrix Analysis and Applications*, Vol. 11, 1990, pp. 335–360.
31. L.N. Trefethen and D. Bau III, *Numerical Linear Algebra*, SIAM, Philadelphia, 1997, pp. 166–170.
32. C.G.J. Jacobi, "Über einen algebraischen Fundamentalsatz und seine Anwendungen [On a Fundamental Theorem in Algebra and Its Applications]," *Journal für die reine und angewandte Mathematik*, Vol. 53, 1857, pp. 275–280.
33. E. Schmidt, "Zur Theorie der linearen und nichtlinearen Integralgleichungen. I Teil, Entwicklung willkürlichen Funktionen nach System vorgeschriebener [On the Theory of Linear and Nonlinear Integral Equations. Part I, Expansion of Arbitrary Functions by a Prescribed System]," *Mathematische Annalen*, Vol. 63, 1907, pp. 433–476.
34. P.S. Laplace, *Theorie Analytique des Probabilités [Analytical Theory of Probabilities]*, 3rd ed., Courcier, Paris, 1820.
35. A.S. Householder, "Unitary Triangularization of a Nonsymmetric Matrix," *J. ACM*, Vol. 5, 1958, pp. 339–342.
36. D. Bogert and W.R. Burris, *Comparison of Least Squares Algorithms*, Tech. Report ORNL-3499, V.1, S5.5, Neutron Physics Div., Oak Ridge Nat'l Laboratory, Oak Ridge, Tenn., 1963.
37. G.H. Golub, "Numerical Methods for Solving Least Squares Problems," *Numerische Mathematik*, Vol. 7, 1965, pp. 206–216.
38. J.J.M. Cuppen, "A Divide and Conquer Method for the Symmetric Eigenproblem," *Numerische Mathematik*, Vol. 36, 1981, pp. 177–195.
39. M. Gu and S.C. Eisenstat, "Stable and Efficient Algorithm for the Rank-One Modification of the Symmetric Eigenproblem," *SIAM J. Matrix Analysis and Applications*, Vol. 15, 1994, pp. 1266–1276.
40. J. Demmel and K. Veselic, "Jacobi's Method Is More Accurate than QR," *SIAM J. Matrix Analysis and Applications*, Vol. 13, 1992, pp. 1204–1245.
41. A.L. Cauchy, "Sur l'equation á l'aide de laquelle on détermine les inégalités séculaires des mouvements des planètes [On the Equation by which the Inequalities for the Secular Movement of Planets Is Determined]," *Oeuvres Completes (II, e Série)*, Vol. 9, 1829.
42. C.G.J. Jacobi, "Über ein leichtes Verfahren die in der Theorie der Säcularstörungen vorkommenden Gleichungen numerisch aufzulösen [On an Easy Method for the Numerical Solution of the Equations that Appear in the Theory of Secular Perturbations]," *Journal für die reine und angewandte Mathematik*, Vol. 30, 1846, pp. 51–94.
43. J.H. Wilkinson, "Householder's Method for the Solution of the Algebraic Eigenvalue Problem," *Computer J.*, Vol. 3, 1960, pp. 23–27.
44. R.H. Bartels and G.W. Stewart, "Algorithm 432: The Solution of the Matrix Equation  $AX - BX = C$ ," *Comm. ACM*, Vol. 8, 1972, pp. 820–826.
45. G.H. Golub, S. Nash, and C. Van Loan, "Hessenberg-Schur Method for the Problem  $AX + XB = C$ ," *IEEE Trans. Automatic Control*, Vol. AC-24, 1979, pp. 909–913.
46. J.G.F. Francis, "The QR Transformation, Parts I and II," *Computer J.*, Vol. 4, 1961, pp. 265–271, 1962, pp. 332–345.
47. J. Schur, "Über die charakteristischen Wurzeln einer linearen Substitution mit einer Anwendung auf die Theorie der Integralgleichungen [On the Characteristic Roots of a Linear Transformation with an Application to the Theory of Integral Equations]," *Mathematische Annalen*, Vol. 66, 1909, pp. 448–510.
49. M. Gu and S.C. Eisenstat, "A Divide-and-Conquer Algorithm for the Bidiagonal SVD," *SIAM J. Matrix Analysis and Applications*, Vol. 16, 1995, pp. 79–92.
50. E. Beltrami, "Sulle funzioni bilineari [On Bilinear Functions]," *Giornale di Matematiche ad Uso degli Studenti Delle Università*, Vol. 11, 1873, pp. 98–106. An English translation by D. Boley is available in Tech. Report 90–37, Dept. of Computer Science, Univ. of Minnesota, Minneapolis, 1990.
51. C. Jordan, "Memoire sur les formes bilineaires [Memoir on Bilinear Forms]," *Journal de Mathematiques Pures et Appliquees, Deuxieme Serie*, Vol. 19, 1874, pp. 35–54.
52. G.H. Golub and W. Kahan, "Calculating the Singular Values and Pseudo-Inverse of a Matrix," *SIAM J. Numerical Analysis*, Vol. 2, 1965, pp. 205–224.
53. E.G. Kogbetliantz, "Solution of Linear Systems by Diagonalization of Coefficients Matrix," *Quarterly of Applied Math.*, Vol. 13, 1955, pp. 123–132.
54. T.F. Chan, "Rank Revealing QR Factorizations," *Linear Algebra and Its Applications*, Vol. 88/89, 1987, pp. 67–82.
55. G.W. Stewart, "An Updating Algorithm for Subspace Tracking," *IEEE Trans. Signal Processing*, Vol. 40, 1992, pp. 1535–1541.
56. Z. Bai and J.W. Demmel, "Computing the Generalized Singular Value Decomposition," *SIAM J. Scientific and Statistical Computing*, Vol. 14, 1993, pp. 1464–1468.
57. C.B. Moler and G.W. Stewart, "An Algorithm for Generalized Matrix Eigenvalue Problems," *SIAM J. Numerical Analysis*, Vol. 10, 1973, pp. 241–256.

G.W. Stewart is a professor in the Department of Computer Science and in the Institute for Advanced Computer Studies at the University of Maryland. His research interests focus on numerical linear algebra, perturbation theory, and rounding-error analysis. He earned his PhD from the University of Tennessee, where his advisor was A.S. Householder. Contact Stewart at the Dept. of Computer Science, Univ. of Maryland, College Park, MD 20742; [stewart@cs.umd.edu](mailto:stewart@cs.umd.edu).

Scientific  
Programming  
will return in  
the next issue  
of *CiSE*.

