

Orthogonal Bases and the QR Algorithm

by Peter J. Olver
University of Minnesota

1. Orthogonal Bases.

Throughout, we work in the Euclidean vector space $V = \mathbb{R}^n$, the space of column vectors with n real entries. As inner product, we will only use the dot product $\mathbf{v} \cdot \mathbf{w} = \mathbf{v}^T \mathbf{w}$ and corresponding Euclidean norm $\|\mathbf{v}\| = \sqrt{\mathbf{v} \cdot \mathbf{v}}$.

Two vectors $\mathbf{v}, \mathbf{w} \in V$ are called *orthogonal* if their inner product vanishes: $\mathbf{v} \cdot \mathbf{w} = 0$. In the case of vectors in Euclidean space, orthogonality under the dot product means that they meet at a right angle. A particularly important configuration is when V admits a basis consisting of mutually orthogonal elements.

Definition 1.1. A basis $\mathbf{u}_1, \dots, \mathbf{u}_n$ of V is called *orthogonal* if $\mathbf{u}_i \cdot \mathbf{u}_j = 0$ for all $i \neq j$. The basis is called *orthonormal* if, in addition, each vector has unit length: $\|\mathbf{u}_i\| = 1$, for all $i = 1, \dots, n$.

The simplest example of an orthonormal basis is the standard basis

$$\mathbf{e}_1 = \begin{pmatrix} 1 \\ 0 \\ 0 \\ \vdots \\ 0 \\ 0 \end{pmatrix}, \quad \mathbf{e}_2 = \begin{pmatrix} 0 \\ 1 \\ 0 \\ \vdots \\ 0 \\ 0 \end{pmatrix}, \quad \dots \quad \mathbf{e}_n = \begin{pmatrix} 0 \\ 0 \\ 0 \\ \vdots \\ 0 \\ 1 \end{pmatrix}.$$

Orthogonality follows because $\mathbf{e}_i \cdot \mathbf{e}_j = 0$, for $i \neq j$, while $\|\mathbf{e}_i\| = 1$ implies normality.

Since a basis cannot contain the zero vector, there is an easy way to convert an orthogonal basis to an orthonormal basis. Namely, we replace each basis vector with a unit vector pointing in the same direction.

Lemma 1.2. If $\mathbf{v}_1, \dots, \mathbf{v}_n$ is an orthogonal basis of a vector space V , then the normalized vectors $\mathbf{u}_i = \mathbf{v}_i / \|\mathbf{v}_i\|$, $i = 1, \dots, n$, form an orthonormal basis.

Example 1.3. The vectors

$$\mathbf{v}_1 = \begin{pmatrix} 1 \\ 2 \\ -1 \end{pmatrix}, \quad \mathbf{v}_2 = \begin{pmatrix} 0 \\ 1 \\ 2 \end{pmatrix}, \quad \mathbf{v}_3 = \begin{pmatrix} 5 \\ -2 \\ 1 \end{pmatrix},$$

are easily seen to form a basis of \mathbb{R}^3 . Moreover, they are mutually perpendicular, $\mathbf{v}_1 \cdot \mathbf{v}_2 = \mathbf{v}_1 \cdot \mathbf{v}_3 = \mathbf{v}_2 \cdot \mathbf{v}_3 = 0$, and so form an orthogonal basis with respect to the standard dot

product on \mathbb{R}^3 . When we divide each orthogonal basis vector by its length, the result is the orthonormal basis

$$\mathbf{u}_1 = \frac{1}{\sqrt{6}} \begin{pmatrix} 1 \\ 2 \\ -1 \end{pmatrix} = \begin{pmatrix} \frac{1}{\sqrt{6}} \\ \frac{2}{\sqrt{6}} \\ -\frac{1}{\sqrt{6}} \end{pmatrix}, \quad \mathbf{u}_2 = \frac{1}{\sqrt{5}} \begin{pmatrix} 0 \\ 1 \\ 2 \end{pmatrix} = \begin{pmatrix} 0 \\ \frac{1}{\sqrt{5}} \\ \frac{2}{\sqrt{5}} \end{pmatrix}, \quad \mathbf{u}_3 = \frac{1}{\sqrt{30}} \begin{pmatrix} 5 \\ -2 \\ 1 \end{pmatrix} = \begin{pmatrix} \frac{5}{\sqrt{30}} \\ -\frac{2}{\sqrt{30}} \\ \frac{1}{\sqrt{30}} \end{pmatrix},$$

satisfying $\mathbf{u}_1 \cdot \mathbf{u}_2 = \mathbf{u}_1 \cdot \mathbf{u}_3 = \mathbf{u}_2 \cdot \mathbf{u}_3 = 0$ and $\|\mathbf{u}_1\| = \|\mathbf{u}_2\| = \|\mathbf{u}_3\| = 1$. The appearance of square roots in the elements of an orthonormal basis is fairly typical.

A useful observation is that any orthogonal collection of nonzero vectors is automatically linearly independent.

Proposition 1.4. *If $\mathbf{v}_1, \dots, \mathbf{v}_k \in V$ are nonzero, mutually orthogonal elements, so $\mathbf{v}_i \neq \mathbf{0}$ and $\mathbf{v}_i \cdot \mathbf{v}_j = 0$ for all $i \neq j$, then they are linearly independent.*

Proof: Suppose

$$c_1 \mathbf{v}_1 + \dots + c_k \mathbf{v}_k = \mathbf{0}.$$

Let us take the inner product of this equation with any \mathbf{v}_i . Using linearity of the inner product and orthogonality, we compute

$$0 = c_1 \mathbf{v}_1 + \dots + c_k \mathbf{v}_k \cdot \mathbf{v}_i = c_1 \mathbf{v}_1 \cdot \mathbf{v}_i + \dots + c_k \mathbf{v}_k \cdot \mathbf{v}_i = c_i \mathbf{v}_i \cdot \mathbf{v}_i = c_i \|\mathbf{v}_i\|^2.$$

Therefore, provided $\mathbf{v}_i \neq \mathbf{0}$, we conclude that the coefficient $c_i = 0$. Since this holds for all $i = 1, \dots, k$, the linear independence of $\mathbf{v}_1, \dots, \mathbf{v}_k$ follows. *Q.E.D.*

As a direct corollary, we infer that any collection of nonzero orthogonal vectors forms a basis for its span.

Theorem 1.5. *Suppose $\mathbf{v}_1, \dots, \mathbf{v}_n \in V$ are nonzero, mutually orthogonal elements of an inner product space V . Then $\mathbf{v}_1, \dots, \mathbf{v}_n$ form an orthogonal basis for their span $W = \text{span}\{\mathbf{v}_1, \dots, \mathbf{v}_n\} \subset V$, which is therefore a subspace of dimension $n = \dim W$. In particular, if $\dim V = n$, then $\mathbf{v}_1, \dots, \mathbf{v}_n$ form a orthogonal basis for V .*

Computations in Orthogonal Bases

What are the advantages of orthogonal and orthonormal bases? Once one has a basis of a vector space, a key issue is how to express other elements as linear combinations of the basis elements — that is, to find their *coordinates* in the prescribed basis. In general, this is not so easy, since it requires solving a system of linear equations. In high dimensional situations arising in applications, computing the solution may require a considerable, if not infeasible amount of time and effort.

However, if the basis is orthogonal, or, even better, orthonormal, then the change of basis computation requires almost no work. This is the crucial insight underlying the efficacy of both discrete and continuous Fourier analysis, least squares approximations and statistical analysis of large data sets, signal, image and video processing, and a multitude of other applications, both classical and modern.

Theorem 1.6. Let $\mathbf{u}_1, \dots, \mathbf{u}_n$ be an orthonormal basis for an inner product space V . Then one can write any element $\mathbf{v} \in V$ as a linear combination

$$\mathbf{v} = c_1 \mathbf{u}_1 + \cdots + c_n \mathbf{u}_n, \quad (1.1)$$

in which its coordinates

$$c_i = \mathbf{v} \cdot \mathbf{u}_i, \quad i = 1, \dots, n, \quad (1.2)$$

are explicitly given as inner products. Moreover, its norm

$$\|\mathbf{v}\| = \sqrt{c_1^2 + \cdots + c_n^2} = \sqrt{\sum_{i=1}^n \mathbf{v} \cdot \mathbf{u}_i^2} \quad (1.3)$$

is the square root of the sum of the squares of its orthonormal basis coordinates.

Proof: Let us compute the inner product of (1.1) with one of the basis vectors. Using the orthonormality conditions

$$\mathbf{u}_i \cdot \mathbf{u}_j = \begin{cases} 0 & i \neq j, \\ 1 & i = j, \end{cases} \quad (1.4)$$

and bilinearity of the inner product, we find

$$\mathbf{v} \cdot \mathbf{u}_i = \left\langle \sum_{j=1}^n c_j \mathbf{u}_j; \mathbf{u}_i \right\rangle = \sum_{j=1}^n c_j \mathbf{u}_j \cdot \mathbf{u}_i = c_i \|\mathbf{u}_i\|^2 = c_i.$$

To prove formula (1.3), we similarly expand

$$\|\mathbf{v}\|^2 = \mathbf{v} \cdot \mathbf{v} = \sum_{i,j=1}^n c_i c_j \mathbf{u}_i \cdot \mathbf{u}_j = \sum_{i=1}^n c_i^2,$$

again making use of the orthonormality of the basis elements. *Q.E.D.*

Example 1.7. Let us rewrite the vector $\mathbf{v} = (1, 1, 1)^T$ in terms of the orthonormal basis

$$\mathbf{u}_1 = \begin{pmatrix} \frac{1}{\sqrt{6}} \\ \frac{2}{\sqrt{6}} \\ -\frac{1}{\sqrt{6}} \end{pmatrix}, \quad \mathbf{u}_2 = \begin{pmatrix} 0 \\ \frac{1}{\sqrt{5}} \\ \frac{2}{\sqrt{5}} \end{pmatrix}, \quad \mathbf{u}_3 = \begin{pmatrix} \frac{5}{\sqrt{30}} \\ -\frac{2}{\sqrt{30}} \\ \frac{1}{\sqrt{30}} \end{pmatrix},$$

constructed in Example 1.3. Computing the dot products

$$\mathbf{v} \cdot \mathbf{u}_1 = \frac{2}{\sqrt{6}}, \quad \mathbf{v} \cdot \mathbf{u}_2 = \frac{3}{\sqrt{5}}, \quad \mathbf{v} \cdot \mathbf{u}_3 = \frac{4}{\sqrt{30}},$$

we immediately conclude that

$$\mathbf{v} = \frac{2}{\sqrt{6}} \mathbf{u}_1 + \frac{3}{\sqrt{5}} \mathbf{u}_2 + \frac{4}{\sqrt{30}} \mathbf{u}_3.$$

Needless to say, a direct computation based on solving the associated linear system is more tedious.

While passage from an orthogonal basis to its orthonormal version is elementary — one simply divides each basis element by its norm — we shall often find it more convenient to work directly with the unnormalized version. The next result provides the corresponding formula expressing a vector in terms of an orthogonal, but not necessarily orthonormal basis. The proof proceeds exactly as in the orthonormal case, and details are left to the reader.

Theorem 1.8. *If $\mathbf{v}_1, \dots, \mathbf{v}_n$ form an orthogonal basis, then the corresponding coordinates of a vector*

$$\mathbf{v} = a_1 \mathbf{v}_1 + \cdots + a_n \mathbf{v}_n \quad \text{are given by} \quad a_i = \frac{\mathbf{v} \cdot \mathbf{v}_i}{\|\mathbf{v}_i\|^2}. \quad (1.5)$$

In this case, its norm can be computed using the formula

$$\|\mathbf{v}\|^2 = \sum_{i=1}^n a_i^2 \|\mathbf{v}_i\|^2 = \sum_{i=1}^n \left(\frac{\mathbf{v} \cdot \mathbf{v}_i}{\|\mathbf{v}_i\|} \right)^2. \quad (1.6)$$

Equation (1.5), along with its orthonormal simplification (1.2), is one of the most useful formulas we shall establish, and applications will appear repeatedly throughout the sequel.

Example 1.9. The wavelet basis

$$\mathbf{v}_1 = \begin{pmatrix} 1 \\ 1 \\ 1 \\ 1 \end{pmatrix}, \quad \mathbf{v}_2 = \begin{pmatrix} 1 \\ 1 \\ -1 \\ -1 \end{pmatrix}, \quad \mathbf{v}_3 = \begin{pmatrix} 1 \\ -1 \\ 0 \\ 0 \end{pmatrix}, \quad \mathbf{v}_4 = \begin{pmatrix} 0 \\ 0 \\ 1 \\ -1 \end{pmatrix}, \quad (1.7)$$

is an orthogonal basis of \mathbb{R}^4 . The norms are

$$\|\mathbf{v}_1\| = 2, \quad \|\mathbf{v}_2\| = 2, \quad \|\mathbf{v}_3\| = \sqrt{2}, \quad \|\mathbf{v}_4\| = \sqrt{2}.$$

Therefore, using (1.5), we can readily express any vector as a linear combination of the wavelet basis vectors. For example,

$$\mathbf{v} = \begin{pmatrix} 4 \\ -2 \\ 1 \\ 5 \end{pmatrix} = 2 \mathbf{v}_1 - \mathbf{v}_2 + 3 \mathbf{v}_3 - 2 \mathbf{v}_4,$$

where the wavelet coordinates are computed directly by

$$\frac{\mathbf{v} \cdot \mathbf{v}_1}{\|\mathbf{v}_1\|^2} = \frac{8}{4} = 2, \quad \frac{\mathbf{v} \cdot \mathbf{v}_2}{\|\mathbf{v}_2\|^2} = \frac{-4}{4} = -1, \quad \frac{\mathbf{v} \cdot \mathbf{v}_3}{\|\mathbf{v}_3\|^2} = \frac{6}{2} = 3, \quad \frac{\mathbf{v} \cdot \mathbf{v}_4}{\|\mathbf{v}_4\|^2} = \frac{-4}{2} = -2.$$

Finally, we note that

$$46 = \|\mathbf{v}\|^2 = 2^2 \|\mathbf{v}_1\|^2 + (-1)^2 \|\mathbf{v}_2\|^2 + 3^2 \|\mathbf{v}_3\|^2 + (-2)^2 \|\mathbf{v}_4\|^2 = 4 \cdot 4 + 1 \cdot 4 + 9 \cdot 2 + 4 \cdot 2,$$

in conformity with (1.6).

2. The Gram–Schmidt Process.

Once we become convinced of the utility of orthogonal and orthonormal bases, a natural question arises: How can we construct them? A practical algorithm was first discovered by Pierre–Simon Laplace in the eighteenth century. Today the algorithm is known as the *Gram–Schmidt process*, after its rediscovery by the nineteenth century mathematicians Jorgen Gram and Erhard Schmidt. The Gram–Schmidt process is one of the premier algorithms of applied and computational linear algebra.

We assume that we already know some basis $\mathbf{w}_1, \dots, \mathbf{w}_n$ of V , where $n = \dim V$. Our goal is to use this information to construct an orthogonal basis $\mathbf{v}_1, \dots, \mathbf{v}_n$. We will construct the orthogonal basis elements one by one. Since initially we are not worrying about normality, there are no conditions on the first orthogonal basis element \mathbf{v}_1 , and so there is no harm in choosing

$$\mathbf{v}_1 = \mathbf{w}_1.$$

Note that $\mathbf{v}_1 \neq \mathbf{0}$, since \mathbf{w}_1 appears in the original basis. The second basis vector must be orthogonal to the first: $\mathbf{v}_2 \cdot \mathbf{v}_1 = 0$. Let us try to arrange this by subtracting a suitable multiple of \mathbf{v}_1 , and set

$$\mathbf{v}_2 = \mathbf{w}_2 - c\mathbf{v}_1,$$

where c is a scalar to be determined. The orthogonality condition

$$0 = \mathbf{v}_2 \cdot \mathbf{v}_1 = \mathbf{w}_2 \cdot \mathbf{v}_1 - c\mathbf{v}_1 \cdot \mathbf{v}_1 = \mathbf{w}_2 \cdot \mathbf{v}_1 - c\|\mathbf{v}_1\|^2$$

requires that $c = \mathbf{w}_2 \cdot \mathbf{v}_1 / \|\mathbf{v}_1\|^2$, and therefore

$$\mathbf{v}_2 = \mathbf{w}_2 - \frac{\mathbf{w}_2 \cdot \mathbf{v}_1}{\|\mathbf{v}_1\|^2} \mathbf{v}_1. \quad (2.1)$$

Linear independence of $\mathbf{v}_1 = \mathbf{w}_1$ and \mathbf{w}_2 ensures that $\mathbf{v}_2 \neq \mathbf{0}$. (Check!)

Next, we construct

$$\mathbf{v}_3 = \mathbf{w}_3 - c_1\mathbf{v}_1 - c_2\mathbf{v}_2$$

by subtracting suitable multiples of the first two orthogonal basis elements from \mathbf{w}_3 . We want \mathbf{v}_3 to be orthogonal to both \mathbf{v}_1 and \mathbf{v}_2 . Since we already arranged that $\mathbf{v}_1 \cdot \mathbf{v}_2 = 0$, this requires

$$0 = \mathbf{v}_3 \cdot \mathbf{v}_1 = \mathbf{w}_3 \cdot \mathbf{v}_1 - c_1\mathbf{v}_1 \cdot \mathbf{v}_1, \quad 0 = \mathbf{v}_3 \cdot \mathbf{v}_2 = \mathbf{w}_3 \cdot \mathbf{v}_2 - c_2\mathbf{v}_2 \cdot \mathbf{v}_2,$$

and hence

$$c_1 = \frac{\mathbf{w}_3 \cdot \mathbf{v}_1}{\|\mathbf{v}_1\|^2}, \quad c_2 = \frac{\mathbf{w}_3 \cdot \mathbf{v}_2}{\|\mathbf{v}_2\|^2}.$$

Therefore, the next orthogonal basis vector is given by the formula

$$\mathbf{v}_3 = \mathbf{w}_3 - \frac{\mathbf{w}_3 \cdot \mathbf{v}_1}{\|\mathbf{v}_1\|^2} \mathbf{v}_1 - \frac{\mathbf{w}_3 \cdot \mathbf{v}_2}{\|\mathbf{v}_2\|^2} \mathbf{v}_2.$$

Since \mathbf{v}_1 and \mathbf{v}_2 are linear combinations of \mathbf{w}_1 and \mathbf{w}_2 , we must have $\mathbf{v}_3 \neq \mathbf{0}$, as otherwise this would imply that $\mathbf{w}_1, \mathbf{w}_2, \mathbf{w}_3$ are linearly dependent, and hence could not come from a basis.

Continuing in the same manner, suppose we have already constructed the mutually orthogonal vectors $\mathbf{v}_1, \dots, \mathbf{v}_{k-1}$ as linear combinations of $\mathbf{w}_1, \dots, \mathbf{w}_{k-1}$. The next orthogonal basis element \mathbf{v}_k will be obtained from \mathbf{w}_k by subtracting off a suitable linear combination of the previous orthogonal basis elements:

$$\mathbf{v}_k = \mathbf{w}_k - c_1 \mathbf{v}_1 - \cdots - c_{k-1} \mathbf{v}_{k-1}.$$

Since $\mathbf{v}_1, \dots, \mathbf{v}_{k-1}$ are already orthogonal, the orthogonality constraint

$$0 = \mathbf{v}_k \cdot \mathbf{v}_j = \mathbf{w}_k \cdot \mathbf{v}_j - c_j \mathbf{v}_j \cdot \mathbf{v}_j$$

requires

$$c_j = \frac{\mathbf{w}_k \cdot \mathbf{v}_j}{\|\mathbf{v}_j\|^2} \quad \text{for} \quad j = 1, \dots, k-1. \quad (2.2)$$

In this fashion, we establish the general *Gram–Schmidt formula*

$$\mathbf{v}_k = \mathbf{w}_k - \sum_{j=1}^{k-1} \frac{\mathbf{w}_k \cdot \mathbf{v}_j}{\|\mathbf{v}_j\|^2} \mathbf{v}_j, \quad k = 1, \dots, n. \quad (2.3)$$

The Gram–Schmidt process (2.3) defines an explicit, recursive procedure for constructing the orthogonal basis vectors $\mathbf{v}_1, \dots, \mathbf{v}_n$. If we are actually after an orthonormal basis $\mathbf{u}_1, \dots, \mathbf{u}_n$, we merely normalize the resulting orthogonal basis vectors, setting $\mathbf{u}_k = \mathbf{v}_k / \|\mathbf{v}_k\|$ for each $k = 1, \dots, n$.

Example 2.1. The vectors

$$\mathbf{w}_1 = \begin{pmatrix} 1 \\ 1 \\ -1 \end{pmatrix}, \quad \mathbf{w}_2 = \begin{pmatrix} 1 \\ 0 \\ 2 \end{pmatrix}, \quad \mathbf{w}_3 = \begin{pmatrix} 2 \\ -2 \\ 3 \end{pmatrix}, \quad (2.4)$$

are readily seen to form a basis[†] of \mathbb{R}^3 . To construct an orthogonal basis (with respect to the standard dot product) using the Gram–Schmidt procedure, we begin by setting

$$\mathbf{v}_1 = \mathbf{w}_1 = \begin{pmatrix} 1 \\ 1 \\ -1 \end{pmatrix}.$$

The next basis vector is

$$\mathbf{v}_2 = \mathbf{w}_2 - \frac{\mathbf{w}_2 \cdot \mathbf{v}_1}{\|\mathbf{v}_1\|^2} \mathbf{v}_1 = \begin{pmatrix} 1 \\ 0 \\ 2 \end{pmatrix} - \frac{-1}{3} \begin{pmatrix} 1 \\ 1 \\ -1 \end{pmatrix} = \begin{pmatrix} \frac{4}{3} \\ \frac{1}{3} \\ \frac{5}{3} \end{pmatrix}.$$

[†] This will, in fact, be a consequence of the successful completion of the Gram–Schmidt process and does not need to be checked in advance. If the given vectors were not linearly independent, then eventually one of the Gram–Schmidt vectors would vanish, $\mathbf{v}_k = \mathbf{0}$, and the process will break down.

The last orthogonal basis vector is

$$\mathbf{v}_3 = \mathbf{w}_3 - \frac{\mathbf{w}_3 \cdot \mathbf{v}_1}{\|\mathbf{v}_1\|^2} \mathbf{v}_1 - \frac{\mathbf{w}_3 \cdot \mathbf{v}_2}{\|\mathbf{v}_2\|^2} \mathbf{v}_2 = \begin{pmatrix} 2 \\ -2 \\ 3 \end{pmatrix} - \frac{-3}{3} \begin{pmatrix} 1 \\ 1 \\ -1 \end{pmatrix} - \frac{7}{\frac{14}{3}} \begin{pmatrix} \frac{4}{3} \\ \frac{1}{3} \\ \frac{5}{3} \end{pmatrix} = \begin{pmatrix} 1 \\ -\frac{3}{2} \\ -\frac{1}{2} \end{pmatrix}.$$

The reader can easily validate the orthogonality of $\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3$.

An orthonormal basis is obtained by dividing each vector by its length. Since

$$\|\mathbf{v}_1\| = \sqrt{3}, \quad \|\mathbf{v}_2\| = \sqrt{\frac{14}{3}}, \quad \|\mathbf{v}_3\| = \sqrt{\frac{7}{2}}.$$

we produce the corresponding orthonormal basis vectors

$$\mathbf{u}_1 = \begin{pmatrix} \frac{1}{\sqrt{3}} \\ \frac{1}{\sqrt{3}} \\ -\frac{1}{\sqrt{3}} \end{pmatrix}, \quad \mathbf{u}_2 = \begin{pmatrix} \frac{4}{\sqrt{42}} \\ \frac{1}{\sqrt{42}} \\ \frac{5}{\sqrt{42}} \end{pmatrix}, \quad \mathbf{u}_3 = \begin{pmatrix} \frac{2}{\sqrt{14}} \\ -\frac{3}{\sqrt{14}} \\ -\frac{1}{\sqrt{14}} \end{pmatrix}. \quad (2.5)$$

Modifications of the Gram–Schmidt Process

With the basic Gram–Schmidt algorithm now in hand, it is worth looking at a couple of reformulations that have both practical and theoretical advantages. The first can be used to directly construct the orthonormal basis vectors $\mathbf{u}_1, \dots, \mathbf{u}_n$ from the basis $\mathbf{w}_1, \dots, \mathbf{w}_n$.

We begin by replacing each orthogonal basis vector in the basic Gram–Schmidt formula (2.3) by its normalized version $\mathbf{u}_j = \mathbf{v}_j / \|\mathbf{v}_j\|$. The original basis vectors can be expressed in terms of the orthonormal basis via a “triangular” system

$$\begin{aligned} \mathbf{w}_1 &= r_{11} \mathbf{u}_1, \\ \mathbf{w}_2 &= r_{12} \mathbf{u}_1 + r_{22} \mathbf{u}_2, \\ \mathbf{w}_3 &= r_{13} \mathbf{u}_1 + r_{23} \mathbf{u}_2 + r_{33} \mathbf{u}_3, \\ &\vdots \quad \quad \quad \vdots \quad \quad \quad \vdots \quad \quad \quad \ddots \\ \mathbf{w}_n &= r_{1n} \mathbf{u}_1 + r_{2n} \mathbf{u}_2 + \cdots + r_{nn} \mathbf{u}_n. \end{aligned} \quad (2.6)$$

The coefficients r_{ij} can, in fact, be computed directly from these formulae. Indeed, taking the inner product of the equation for \mathbf{w}_j with the orthonormal basis vector \mathbf{u}_i for $i \leq j$, we find, in view of the orthonormality constraints (1.4),

$$\mathbf{w}_j \cdot \mathbf{u}_i = r_{1j} \mathbf{u}_1 \cdot \mathbf{u}_i + \cdots + r_{jj} \mathbf{u}_j \cdot \mathbf{u}_i = r_{1j} \mathbf{u}_1 \cdot \mathbf{u}_i + \cdots + r_{jj} \mathbf{u}_n \cdot \mathbf{u}_i = r_{ij},$$

and hence

$$r_{ij} = \mathbf{w}_j \cdot \mathbf{u}_i. \quad (2.7)$$

On the other hand, according to (1.3),

$$\|\mathbf{w}_j\|^2 = \|r_{1j} \mathbf{u}_1 + \cdots + r_{jj} \mathbf{u}_j\|^2 = r_{1j}^2 + \cdots + r_{j-1,j}^2 + r_{jj}^2. \quad (2.8)$$

The pair of equations (2.7–8) can be rearranged to devise a recursive procedure to compute the orthonormal basis. At stage j , we assume that we have already constructed $\mathbf{u}_1, \dots, \mathbf{u}_{j-1}$. We then compute[†]

$$r_{ij} = \mathbf{w}_j \cdot \mathbf{u}_i, \quad \text{for each } i = 1, \dots, j-1. \quad (2.9)$$

We obtain the next orthonormal basis vector \mathbf{u}_j by computing

$$r_{jj} = \sqrt{\|\mathbf{w}_j\|^2 - r_{1j}^2 - \dots - r_{j-1,j}^2}, \quad \mathbf{u}_j = \frac{\mathbf{w}_j - r_{1j}\mathbf{u}_1 - \dots - r_{j-1,j}\mathbf{u}_{j-1}}{r_{jj}}. \quad (2.10)$$

Running through the formulae (2.9–10) for $j = 1, \dots, n$ leads to the *same* orthonormal basis $\mathbf{u}_1, \dots, \mathbf{u}_n$ as the previous version of the Gram–Schmidt procedure.

Example 2.2. Let us apply the revised algorithm to the vectors

$$\mathbf{w}_1 = \begin{pmatrix} 1 \\ 1 \\ -1 \end{pmatrix}, \quad \mathbf{w}_2 = \begin{pmatrix} 1 \\ 0 \\ 2 \end{pmatrix}, \quad \mathbf{w}_3 = \begin{pmatrix} 2 \\ -2 \\ 3 \end{pmatrix},$$

of Example 2.1. To begin, we set

$$r_{11} = \|\mathbf{w}_1\| = \sqrt{3}, \quad \mathbf{u}_1 = \frac{\mathbf{w}_1}{r_{11}} = \begin{pmatrix} \frac{1}{\sqrt{3}} \\ \frac{1}{\sqrt{3}} \\ -\frac{1}{\sqrt{3}} \end{pmatrix}.$$

The next step is to compute

$$r_{12} = \mathbf{w}_2 \cdot \mathbf{u}_1 = -\frac{1}{\sqrt{3}}, \quad r_{22} = \sqrt{\|\mathbf{w}_2\|^2 - r_{12}^2} = \sqrt{\frac{14}{3}}, \quad \mathbf{u}_2 = \frac{\mathbf{w}_2 - r_{12}\mathbf{u}_1}{r_{22}} = \begin{pmatrix} \frac{4}{\sqrt{42}} \\ \frac{1}{\sqrt{42}} \\ \frac{5}{\sqrt{42}} \end{pmatrix}.$$

The final step yields

$$r_{13} = \mathbf{w}_3 \cdot \mathbf{u}_1 = -\sqrt{3}, \quad r_{23} = \mathbf{w}_3 \cdot \mathbf{u}_2 = \sqrt{\frac{21}{2}},$$

$$r_{33} = \sqrt{\|\mathbf{w}_3\|^2 - r_{13}^2 - r_{23}^2} = \sqrt{\frac{7}{2}}, \quad \mathbf{u}_3 = \frac{\mathbf{w}_3 - r_{13}\mathbf{u}_1 - r_{23}\mathbf{u}_2}{r_{33}} = \begin{pmatrix} \frac{2}{\sqrt{14}} \\ -\frac{3}{\sqrt{14}} \\ -\frac{1}{\sqrt{14}} \end{pmatrix}.$$

As advertised, the result is the same orthonormal basis vectors that we found in Example 2.1.

[†] When $j = 1$, there is nothing to do.

For hand computations, the original version (2.3) of the Gram–Schmidt process is slightly easier — even if one does ultimately want an orthonormal basis — since it avoids the square roots that are ubiquitous in the orthonormal version (2.9–10). On the other hand, for numerical implementation on a computer, the orthonormal version is a bit faster, as it involves fewer arithmetic operations.

However, in practical, large scale computations, both versions of the Gram–Schmidt process suffer from a serious flaw. They are subject to numerical instabilities, and so accumulating round-off errors may seriously corrupt the computations, leading to inaccurate, non-orthogonal vectors. Fortunately, there is a simple rearrangement of the calculation that obviates this difficulty and leads to the numerically robust algorithm that is most often used in practice. The idea is to treat the vectors simultaneously rather than sequentially, making full use of the orthonormal basis vectors as they arise. More specifically, the algorithm begins as before — we take $\mathbf{u}_1 = \mathbf{w}_1 / \|\mathbf{w}_1\|$. We then subtract off the appropriate multiples of \mathbf{u}_1 from *all* of the remaining basis vectors so as to arrange their orthogonality to \mathbf{u}_1 . This is accomplished by setting

$$\mathbf{w}_k^{(2)} = \mathbf{w}_k - \mathbf{w}_k \cdot \mathbf{u}_1 \mathbf{u}_1 \quad \text{for} \quad k = 2, \dots, n.$$

The second orthonormal basis vector $\mathbf{u}_2 = \mathbf{w}_2^{(2)} / \|\mathbf{w}_2^{(2)}\|$ is then obtained by normalizing. We next modify the remaining $\mathbf{w}_3^{(2)}, \dots, \mathbf{w}_n^{(2)}$ to produce vectors

$$\mathbf{w}_k^{(3)} = \mathbf{w}_k^{(2)} - \mathbf{w}_k^{(2)} \cdot \mathbf{u}_2 \mathbf{u}_2, \quad k = 3, \dots, n,$$

that are orthogonal to both \mathbf{u}_1 and \mathbf{u}_2 . Then $\mathbf{u}_3 = \mathbf{w}_3^{(3)} / \|\mathbf{w}_3^{(3)}\|$ is the next orthonormal basis element, and the process continues. The full algorithm starts with the initial basis vectors $\mathbf{w}_j = \mathbf{w}_j^{(1)}$, $j = 1, \dots, n$, and then recursively computes

$$\mathbf{u}_j = \frac{\mathbf{w}_j^{(j)}}{\|\mathbf{w}_j^{(j)}\|}, \quad \mathbf{w}_k^{(j+1)} = \mathbf{w}_k^{(j)} - \mathbf{w}_k^{(j)} \cdot \mathbf{u}_j \mathbf{u}_j, \quad \begin{array}{l} j = 1, \dots, n, \\ k = j + 1, \dots, n. \end{array} \quad (2.11)$$

(In the final phase, when $j = n$, the second formula is no longer needed.) The result is a numerically stable computation of the *same* orthonormal basis vectors $\mathbf{u}_1, \dots, \mathbf{u}_n$.

Example 2.3. Let us apply the stable Gram–Schmidt process (2.11) to the basis vectors

$$\mathbf{w}_1^{(1)} = \mathbf{w}_1 = \begin{pmatrix} 2 \\ 2 \\ -1 \end{pmatrix}, \quad \mathbf{w}_2^{(1)} = \mathbf{w}_2 = \begin{pmatrix} 0 \\ 4 \\ -1 \end{pmatrix}, \quad \mathbf{w}_3^{(1)} = \mathbf{w}_3 = \begin{pmatrix} 1 \\ 2 \\ -3 \end{pmatrix}.$$

The first orthonormal basis vector is $\mathbf{u}_1 = \frac{\mathbf{w}_1^{(1)}}{\|\mathbf{w}_1^{(1)}\|} = \begin{pmatrix} \frac{2}{3} \\ \frac{2}{3} \\ -\frac{1}{3} \end{pmatrix}$. Next, we compute

$$\mathbf{w}_2^{(2)} = \mathbf{w}_2^{(1)} - \mathbf{w}_2^{(1)} \cdot \mathbf{u}_1 \mathbf{u}_1 = \begin{pmatrix} -2 \\ 2 \\ 0 \end{pmatrix}, \quad \mathbf{w}_3^{(2)} = \mathbf{w}_3^{(1)} - \mathbf{w}_3^{(1)} \cdot \mathbf{u}_1 \mathbf{u}_1 = \begin{pmatrix} -1 \\ 0 \\ -2 \end{pmatrix}.$$

The second orthonormal basis vector is $\mathbf{u}_2 = \frac{\mathbf{w}_2^{(2)}}{\|\mathbf{w}_2^{(2)}\|} = \begin{pmatrix} -\frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} \\ 0 \end{pmatrix}$. Finally,

$$\mathbf{w}_3^{(3)} = \mathbf{w}_3^{(2)} - \mathbf{w}_3^{(2)} \cdot \mathbf{u}_2 \mathbf{u}_2 = \begin{pmatrix} -\frac{1}{2} \\ -\frac{1}{2} \\ -2 \end{pmatrix}, \quad \mathbf{u}_3 = \frac{\mathbf{w}_3^{(3)}}{\|\mathbf{w}_3^{(3)}\|} = \begin{pmatrix} -\frac{\sqrt{2}}{6} \\ -\frac{\sqrt{2}}{6} \\ -\frac{2\sqrt{2}}{3} \end{pmatrix}.$$

The resulting vectors $\mathbf{u}_1, \mathbf{u}_2, \mathbf{u}_3$ form the desired orthonormal basis.

3. Orthogonal Matrices.

Matrices whose columns form an orthonormal basis of \mathbb{R}^n relative to the standard Euclidean dot product play a distinguished role. Such “orthogonal matrices” appear in a wide range of applications in geometry, physics, quantum mechanics, crystallography, partial differential equations, symmetry theory, and special functions. Rotational motions of bodies in three-dimensional space are described by orthogonal matrices, and hence they lie at the foundations of rigid body mechanics, including satellite and underwater vehicle motions, as well as three-dimensional computer graphics and animation. Furthermore, orthogonal matrices are an essential ingredient in one of the most important methods of numerical linear algebra: the QR algorithm for computing eigenvalues of matrices.

Definition 3.1. A square matrix Q is called an *orthogonal matrix* if it satisfies

$$Q^T Q = I. \quad (3.1)$$

The orthogonality condition implies that one can easily invert an orthogonal matrix:

$$Q^{-1} = Q^T. \quad (3.2)$$

In fact, the two conditions are equivalent, and hence a matrix is orthogonal if and only if its inverse is equal to its transpose. The second important characterization of orthogonal matrices relates them directly to orthonormal bases.

Proposition 3.2. A matrix Q is orthogonal if and only if its columns form an orthonormal basis with respect to the Euclidean dot product on \mathbb{R}^n .

Proof: Let $\mathbf{u}_1, \dots, \mathbf{u}_n$ be the columns of Q . Then $\mathbf{u}_1^T, \dots, \mathbf{u}_n^T$ are the rows of the transposed matrix Q^T . The (i, j) entry of the product $Q^T Q$ is given as the product of the i^{th} row of Q^T times the j^{th} column of Q . Thus, the orthogonality requirement (3.1) implies $\mathbf{u}_i \cdot \mathbf{u}_j = \mathbf{u}_i^T \mathbf{u}_j = \begin{cases} 1, & i = j, \\ 0, & i \neq j, \end{cases}$ which are precisely the conditions (1.4) for $\mathbf{u}_1, \dots, \mathbf{u}_n$ to form an orthonormal basis. *Q.E.D.*

Warning: Technically, we should be referring to an “orthonormal” matrix, not an “orthogonal” matrix. But the terminology is so standard throughout mathematics that we have no choice but to adopt it here. There is no commonly accepted name for a matrix whose columns form an orthogonal but not orthonormal basis.

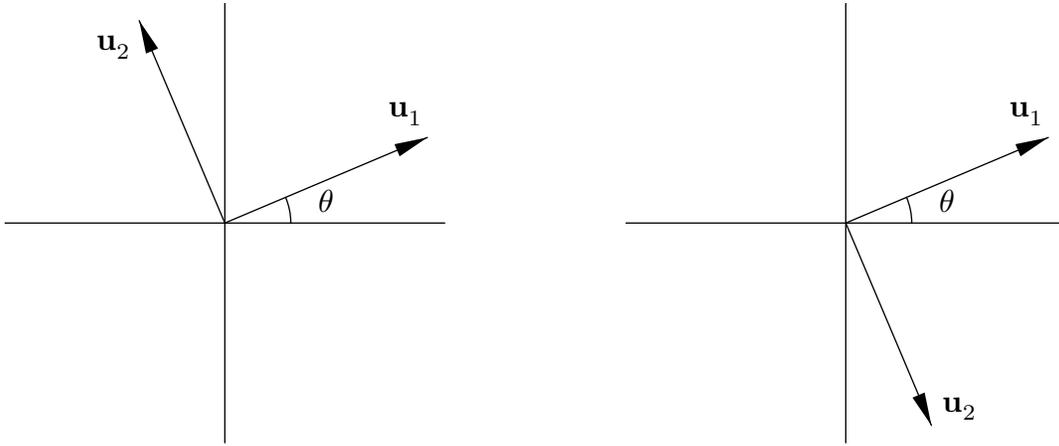


Figure 1. Orthonormal Bases in \mathbb{R}^2 .

Example 3.3. A 2×2 matrix $Q = \begin{pmatrix} a & b \\ c & d \end{pmatrix}$ is orthogonal if and only if its columns $\mathbf{u}_1 = \begin{pmatrix} a \\ c \end{pmatrix}$, $\mathbf{u}_2 = \begin{pmatrix} b \\ d \end{pmatrix}$, form an orthonormal basis of \mathbb{R}^2 . Equivalently, the requirement

$$Q^T Q = \begin{pmatrix} a & c \\ b & d \end{pmatrix} \begin{pmatrix} a & b \\ c & d \end{pmatrix} = \begin{pmatrix} a^2 + c^2 & ab + cd \\ ab + cd & b^2 + d^2 \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix},$$

implies that its entries must satisfy the algebraic equations

$$a^2 + c^2 = 1, \quad ab + cd = 0, \quad b^2 + d^2 = 1.$$

The first and last equations say that the points $(a, c)^T$ and $(b, d)^T$ lie on the unit circle in \mathbb{R}^2 , and so

$$a = \cos \theta, \quad c = \sin \theta, \quad b = \cos \psi, \quad d = \sin \psi,$$

for some choice of angles θ, ψ . The remaining orthogonality condition is

$$0 = ab + cd = \cos \theta \cos \psi + \sin \theta \sin \psi = \cos(\theta - \psi),$$

which implies that θ and ψ differ by a right angle: $\psi = \theta \pm \frac{1}{2}\pi$. The \pm sign leads to two cases:

$$b = -\sin \theta, \quad d = \cos \theta, \quad \text{or} \quad b = \sin \theta, \quad d = -\cos \theta.$$

As a result, every 2×2 orthogonal matrix has one of two possible forms

$$\begin{pmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{pmatrix} \quad \text{or} \quad \begin{pmatrix} \cos \theta & \sin \theta \\ \sin \theta & -\cos \theta \end{pmatrix}, \quad \text{where} \quad 0 \leq \theta < 2\pi. \quad (3.3)$$

The corresponding orthonormal bases are illustrated in Figure 1. The former is a right handed basis, and can be obtained from the standard basis $\mathbf{e}_1, \mathbf{e}_2$ by a rotation through angle θ , while the latter has the opposite, reflected orientation.

Example 3.4. A 3×3 orthogonal matrix $Q = (\mathbf{u}_1 \ \mathbf{u}_2 \ \mathbf{u}_3)$ is prescribed by 3 mutually perpendicular vectors of unit length in \mathbb{R}^3 . For instance, the orthonormal basis constructed

in (2.5) corresponds to the orthogonal matrix $Q = \begin{pmatrix} \frac{1}{\sqrt{3}} & \frac{4}{\sqrt{42}} & \frac{2}{\sqrt{14}} \\ \frac{1}{\sqrt{3}} & \frac{1}{\sqrt{42}} & -\frac{3}{\sqrt{14}} \\ -\frac{1}{\sqrt{3}} & \frac{5}{\sqrt{42}} & -\frac{1}{\sqrt{14}} \end{pmatrix}$.

Lemma 3.5. *An orthogonal matrix has determinant $\det Q = \pm 1$.*

Proof: Taking the determinant of (3.1),

$$1 = \det I = \det(Q^T Q) = \det Q^T \det Q = (\det Q)^2,$$

which immediately proves the lemma.

Q.E.D.

An orthogonal matrix is called *proper* or *special* if it has determinant $+1$. Geometrically, the columns of a proper orthogonal matrix form a right handed basis of \mathbb{R}^n . An *improper* orthogonal matrix, with determinant -1 , corresponds to a left handed basis that lives in a mirror image world.

Proposition 3.6. *The product of two orthogonal matrices is also orthogonal.*

Proof: If $Q_1^T Q_1 = I = Q_2^T Q_2$, then $(Q_1 Q_2)^T (Q_1 Q_2) = Q_2^T Q_1^T Q_1 Q_2 = Q_2^T Q_2 = I$, and so the product matrix $Q_1 Q_2$ is also orthogonal. *Q.E.D.*

This property says that the set of all orthogonal matrices forms a *group*. The *orthogonal group* lies at the foundation of everyday Euclidean geometry, as well as rigid body mechanics, atomic structure and chemistry, computer graphics and animation, and many other areas.

4. Eigenvalues of Symmetric Matrices.

Symmetric matrices play an important role in a broad range of applications, and enjoy a number of important properties not shared by more general matrices. Not only are the eigenvalues of a symmetric matrix necessarily real, the eigenvectors always form an *orthogonal basis*. In fact, this is by far the most common way for orthogonal bases to appear — as the eigenvector bases of symmetric matrices. Let us state this important result, but defer its proof until the end of the section.

Theorem 4.1. *Let $A = A^T$ be a real symmetric $n \times n$ matrix. Then*

- (a) *All the eigenvalues of A are real.*
- (b) *Eigenvectors corresponding to distinct eigenvalues are orthogonal.*
- (c) *There is an orthonormal basis of \mathbb{R}^n consisting of n eigenvectors of A .*

In particular, all symmetric matrices are complete.

Example 4.2. Consider the symmetric matrix $A = \begin{pmatrix} 5 & -4 & 2 \\ -4 & 5 & 2 \\ 2 & 2 & -1 \end{pmatrix}$. A straightforward computation produces its eigenvalues and eigenvectors:

$$\begin{aligned} \lambda_1 &= 9, & \lambda_2 &= 3, & \lambda_3 &= -3, \\ \mathbf{v}_1 &= \begin{pmatrix} 1 \\ -1 \\ 0 \end{pmatrix}, & \mathbf{v}_2 &= \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix}, & \mathbf{v}_3 &= \begin{pmatrix} 1 \\ 1 \\ -2 \end{pmatrix}. \end{aligned}$$

As the reader can check, the eigenvectors form an orthogonal basis of \mathbb{R}^3 . An orthonormal basis is provided by the unit eigenvectors

$$\mathbf{u}_1 = \begin{pmatrix} \frac{1}{\sqrt{2}} \\ -\frac{1}{\sqrt{2}} \\ 0 \end{pmatrix}, \quad \mathbf{u}_2 = \begin{pmatrix} \frac{1}{\sqrt{3}} \\ \frac{1}{\sqrt{3}} \\ \frac{1}{\sqrt{3}} \end{pmatrix}, \quad \mathbf{u}_3 = \begin{pmatrix} \frac{1}{\sqrt{6}} \\ \frac{1}{\sqrt{6}} \\ -\frac{2}{\sqrt{6}} \end{pmatrix}.$$

Proof of Theorem 4.1: First, if $A = A^T$ is real, symmetric, then

$$(A\mathbf{v}) \cdot \mathbf{w} = \mathbf{v} \cdot (A\mathbf{w}) \quad \text{for all } \mathbf{v}, \mathbf{w} \in \mathbb{C}^n, \quad (4.1)$$

where \cdot indicates the Euclidean dot product when the vectors are real and, more generally, the Hermitian dot product $\mathbf{v} \cdot \mathbf{w} = \mathbf{v}^T \overline{\mathbf{w}}$ when they are complex.

To prove property (a), suppose λ is a complex eigenvalue with complex eigenvector $\mathbf{v} \in \mathbb{C}^n$. Consider the Hermitian dot product of the complex vectors $A\mathbf{v}$ and \mathbf{v} :

$$(A\mathbf{v}) \cdot \mathbf{v} = (\lambda\mathbf{v}) \cdot \mathbf{v} = \lambda \|\mathbf{v}\|^2.$$

On the other hand, by (4.1),

$$(A\mathbf{v}) \cdot \mathbf{v} = \mathbf{v} \cdot (A\mathbf{v}) = \mathbf{v} \cdot (\lambda\mathbf{v}) = \mathbf{v}^T \overline{\lambda\mathbf{v}} = \overline{\lambda} \|\mathbf{v}\|^2.$$

Equating these two expressions, we deduce

$$\overline{\lambda} \|\mathbf{v}\|^2 = \lambda \|\mathbf{v}\|^2.$$

Since $\mathbf{v} \neq \mathbf{0}$, as it is an eigenvector, we conclude that $\overline{\lambda} = \lambda$, proving that the eigenvalue λ is real.

To prove (b), suppose

$$A\mathbf{v} = \lambda\mathbf{v}, \quad A\mathbf{w} = \mu\mathbf{w},$$

where $\lambda \neq \mu$ are distinct real eigenvalues. Then, again by (4.1),

$$\lambda\mathbf{v} \cdot \mathbf{w} = (A\mathbf{v}) \cdot \mathbf{w} = \mathbf{v} \cdot (A\mathbf{w}) = \mathbf{v} \cdot (\mu\mathbf{w}) = \mu\mathbf{v} \cdot \mathbf{w},$$

and hence

$$(\lambda - \mu)\mathbf{v} \cdot \mathbf{w} = 0.$$

Since $\lambda \neq \mu$, this implies that $\mathbf{v} \cdot \mathbf{w} = 0$ and hence the eigenvectors \mathbf{v}, \mathbf{w} are orthogonal.

Finally, the proof of (c) is easy if all the eigenvalues of A are distinct. Part (b) proves that the corresponding eigenvectors are orthogonal, and Proposition 1.4 proves that they form a basis. To obtain an orthonormal basis, we merely divide the eigenvectors by their lengths: $\mathbf{u}_k = \mathbf{v}_k / \|\mathbf{v}_k\|$, as in Lemma 1.2.

To prove (c) in general, we proceed by induction on the size n of the matrix A . To start, the case of a 1×1 matrix is trivial. (Why?) Next, suppose A has size $n \times n$. We know that A has at least one eigenvalue, λ_1 , which is necessarily real. Let \mathbf{v}_1 be the associated eigenvector. Let

$$V^\perp = \{ \mathbf{w} \in \mathbb{R}^n \mid \mathbf{v}_1 \cdot \mathbf{w} = 0 \}$$

denote the orthogonal complement to the eigenspace V_{λ_1} — the set of all vectors orthogonal to the first eigenvector. Since $\dim V^\perp = n - 1$, we can choose an orthonormal basis $\mathbf{y}_1, \dots, \mathbf{y}_{n-1}$. Now, if \mathbf{w} is any vector in V^\perp , so is $A\mathbf{w}$, since, by (4.1),

$$\mathbf{v}_1 \cdot (A\mathbf{w}) = (A\mathbf{v}_1) \cdot \mathbf{w} = \lambda_1 \mathbf{v}_1 \cdot \mathbf{w} = 0.$$

Thus, A defines a linear transformation on V^\perp represented by an $(n - 1) \times (n - 1)$ matrix with respect to the chosen orthonormal basis $\mathbf{y}_1, \dots, \mathbf{y}_{n-1}$. It is not hard to prove that the representing matrix is symmetric, and so our induction hypothesis then implies that there is an orthonormal basis of V^\perp consisting of eigenvectors $\mathbf{u}_2, \dots, \mathbf{u}_n$ of A . Appending the unit eigenvector $\mathbf{u}_1 = \mathbf{v}_1 / \|\mathbf{v}_1\|$ to this collection will complete the orthonormal basis of \mathbb{R}^n . *Q.E.D.*

The orthonormal eigenvector basis serves to diagonalize the symmetric matrix, resulting in the so-called *spectral factorization* formula.

Theorem 4.3. *Let A be a real, symmetric matrix. Then there exists an orthogonal matrix Q such that*

$$A = Q \Lambda Q^{-1} = Q \Lambda Q^T, \tag{4.2}$$

where $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_n)$ is a real diagonal matrix. The eigenvalues of A appear on the diagonal of Λ , while the columns of Q are the corresponding orthonormal eigenvectors.

Proof: Equation (4.2) can be rewritten as $AQ = Q\Lambda$. The k^{th} column of the latter matrix equation reads $A\mathbf{v}_k = \lambda_k \mathbf{v}_k$, where \mathbf{v}_k is the k^{th} column of Q . But this is merely the condition that \mathbf{v}_k be an eigenvector of A with eigenvalue λ_k . Theorem 4.1 serves to complete the proof. *Q.E.D.*

Remark: The term “spectrum” refers to the eigenvalues of a matrix or, more generally, a linear operator. The terminology is motivated by physics. The spectral energy lines of atoms, molecules and nuclei are characterized as the eigenvalues of the governing quantum mechanical Schrödinger operator. The *Spectral Theorem* 4.3 is the finite-dimensional version for the decomposition of quantum mechanical linear operators into their spectral eigenstates.

The most important subclass of symmetric matrices are the positive definite matrices.

Definition 4.4. An $n \times n$ symmetric matrix A is called *positive definite* if it satisfies the positivity condition

$$\mathbf{x}^T A \mathbf{x} > 0 \quad \text{for all} \quad \mathbf{0} \neq \mathbf{x} \in \mathbb{R}^n. \quad (4.3)$$

Theorem 4.5. A symmetric matrix $A = A^T$ is positive definite if and only if all of its eigenvalues are strictly positive.

Proof: First, if A is positive definite, and $\mathbf{v} \neq \mathbf{0}$ is an eigenvector with (necessarily real) eigenvalue λ , then (4.3) implies

$$0 < \mathbf{v}^T A \mathbf{v} = \mathbf{v}^T (\lambda \mathbf{v}) = \lambda \|\mathbf{v}\|^2, \quad (4.4)$$

which immediately proves that $\lambda > 0$. Conversely, suppose A has all positive eigenvalues. Let $\mathbf{u}_1, \dots, \mathbf{u}_n$ be the orthonormal eigenvector basis of \mathbb{R}^n guaranteed by Theorem 4.1, with $A \mathbf{u}_j = \lambda_j \mathbf{u}_j$. Then, writing

$$\mathbf{x} = c_1 \mathbf{u}_1 + \dots + c_n \mathbf{u}_n, \quad \text{we find} \quad A \mathbf{x} = c_1 \lambda_1 \mathbf{u}_1 + \dots + c_n \lambda_n \mathbf{u}_n.$$

Therefore, using the orthonormality of the eigenvectors, for any $\mathbf{x} \neq \mathbf{0}$,

$$\mathbf{x}^T A \mathbf{x} = (c_1 \mathbf{u}_1^T + \dots + c_n \mathbf{u}_n^T) (c_1 \lambda_1 \mathbf{u}_1 + \dots + c_n \lambda_n \mathbf{u}_n) = \lambda_1 c_1^2 + \dots + \lambda_n c_n^2 > 0$$

since only $\mathbf{x} = \mathbf{0}$ has $c_1 = \dots = c_n = 0$. This proves that A is positive definite. *Q.E.D.*

5. The QR Factorization.

The Gram–Schmidt procedure for orthonormalizing bases of \mathbb{R}^n can be reinterpreted as a matrix factorization. This is more subtle than the LU factorization that results from Gaussian Elimination, but is of comparable significance, and is used in a broad range of applications in mathematics, statistics, physics, engineering, and numerical analysis.

Let $\mathbf{w}_1, \dots, \mathbf{w}_n$ be a basis of \mathbb{R}^n , and let $\mathbf{u}_1, \dots, \mathbf{u}_n$ be the corresponding orthonormal basis that results from any one of the three implementations of the Gram–Schmidt process. We assemble both sets of column vectors to form nonsingular $n \times n$ matrices

$$A = (\mathbf{w}_1 \quad \mathbf{w}_2 \quad \dots \quad \mathbf{w}_n), \quad Q = (\mathbf{u}_1 \quad \mathbf{u}_2 \quad \dots \quad \mathbf{u}_n).$$

Since the \mathbf{u}_i form an orthonormal basis, Q is an orthogonal matrix. Moreover, the Gram–Schmidt equations (2.6) can be recast into an equivalent matrix form:

$$A = QR, \quad \text{where} \quad R = \begin{pmatrix} r_{11} & r_{12} & \dots & r_{1n} \\ 0 & r_{22} & \dots & r_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & r_{nn} \end{pmatrix} \quad (5.1)$$

is an upper triangular matrix whose entries are the coefficients in (2.9–10). Since the Gram–Schmidt process works on any basis, the only requirement on the matrix A is that its columns form a basis of \mathbb{R}^n , and hence A can be any nonsingular matrix. We have therefore established the celebrated QR factorization of nonsingular matrices.

QR Factorization of a Matrix A

```
start
  for  $j = 1$  to  $n$ 
    set  $r_{jj} = \sqrt{a_{1j}^2 + \dots + a_{nj}^2}$ 
    if  $r_{jj} = 0$ , stop; print “A has linearly dependent columns”
    else for  $i = 1$  to  $n$ 
      set  $a_{ij} = a_{ij}/r_{jj}$ 
    next  $i$ 
    for  $k = j + 1$  to  $n$ 
      set  $r_{jk} = a_{1j}a_{1k} + \dots + a_{nj}a_{nk}$ 
      for  $i = 1$  to  $n$ 
        set  $a_{ik} = a_{ik} - a_{ij}r_{jk}$ 
      next  $i$ 
    next  $k$ 
  next  $j$ 
end
```

Theorem 5.1. Any nonsingular matrix A can be factored, $A = QR$, into the product of an orthogonal matrix Q and an upper triangular matrix R . The factorization is unique if all the diagonal entries of R are assumed to be positive.

We will use the compacted term *positive upper triangular* to refer to upper triangular matrices with positive entries along the diagonal.

Example 5.2. The columns of the matrix $A = \begin{pmatrix} 1 & 1 & 2 \\ 1 & 0 & -2 \\ -1 & 2 & 3 \end{pmatrix}$ are the same as the basis vectors considered in Example 2.2. The orthonormal basis (2.5) constructed using the Gram–Schmidt algorithm leads to the orthogonal and positive upper triangular matrices

$$Q = \begin{pmatrix} \frac{1}{\sqrt{3}} & \frac{4}{\sqrt{42}} & \frac{2}{\sqrt{14}} \\ \frac{1}{\sqrt{3}} & \frac{1}{\sqrt{42}} & -\frac{3}{\sqrt{14}} \\ -\frac{1}{\sqrt{3}} & \frac{5}{\sqrt{42}} & -\frac{1}{\sqrt{14}} \end{pmatrix}, \quad R = \begin{pmatrix} \sqrt{3} & -\frac{1}{\sqrt{3}} & -\sqrt{3} \\ 0 & \frac{\sqrt{14}}{\sqrt{3}} & \frac{\sqrt{21}}{\sqrt{2}} \\ 0 & 0 & \frac{\sqrt{7}}{\sqrt{2}} \end{pmatrix}. \quad (5.2)$$

The reader may wish to verify that, indeed, $A = QR$.

While any of the three implementations of the Gram–Schmidt algorithm will produce the QR factorization of a given matrix $A = (\mathbf{w}_1 \ \mathbf{w}_2 \ \dots \ \mathbf{w}_n)$, the stable version, as encoded in equations (2.11), is the one to use in practical computations, as it is the least likely to fail

due to numerical artifacts produced by round-off errors. The accompanying pseudocode program reformulates the algorithm purely in terms of the matrix entries a_{ij} of A . During the course of the algorithm, the entries of the matrix A are successively overwritten; the final result is the orthogonal matrix Q appearing in place of A . The entries r_{ij} of R must be stored separately.

Example 5.3. Let us factor the matrix

$$A = \begin{pmatrix} 2 & 1 & 0 & 0 \\ 1 & 2 & 1 & 0 \\ 0 & 1 & 2 & 1 \\ 0 & 0 & 1 & 2 \end{pmatrix}$$

using the numerically stable QR algorithm. As in the program, we work directly on the matrix A , gradually changing it into orthogonal form. In the first loop, we set $r_{11} = \sqrt{5}$ to be the norm of the first column vector of A . We then normalize the first column

by dividing by r_{11} ; the resulting matrix is $\begin{pmatrix} \frac{2}{\sqrt{5}} & 1 & 0 & 0 \\ \frac{1}{\sqrt{5}} & 2 & 1 & 0 \\ 0 & 1 & 2 & 1 \\ 0 & 0 & 1 & 2 \end{pmatrix}$. The next entries $r_{12} =$

$\frac{4}{\sqrt{5}}$, $r_{13} = \frac{1}{\sqrt{5}}$, $r_{14} = 0$, are obtained by taking the dot products of the first column with the other three columns. For $j = 1, 2, 3$, we subtract r_{1j} times the first column

from the j^{th} column; the result $\begin{pmatrix} \frac{2}{\sqrt{5}} & -\frac{3}{5} & -\frac{2}{5} & 0 \\ \frac{1}{\sqrt{5}} & \frac{6}{5} & \frac{4}{5} & 0 \\ 0 & 1 & 2 & 1 \\ 0 & 0 & 1 & 2 \end{pmatrix}$ is a matrix whose first column is

normalized to have unit length, and whose second, third and fourth columns are orthogonal to it. In the next loop, we normalize the second column by dividing by its norm $r_{22} =$

$\sqrt{\frac{14}{5}}$, and so obtain the matrix $\begin{pmatrix} \frac{2}{\sqrt{5}} & -\frac{3}{\sqrt{70}} & -\frac{2}{5} & 0 \\ \frac{1}{\sqrt{5}} & \frac{6}{\sqrt{70}} & \frac{4}{5} & 0 \\ 0 & \frac{5}{\sqrt{70}} & 2 & 1 \\ 0 & 0 & 1 & 2 \end{pmatrix}$. We then take dot products of

the second column with the remaining two columns to produce $r_{23} = \frac{16}{\sqrt{70}}$, $r_{24} = \frac{5}{\sqrt{70}}$. Subtracting these multiples of the second column from the third and fourth columns, we

obtain $\begin{pmatrix} \frac{2}{\sqrt{5}} & -\frac{3}{\sqrt{70}} & \frac{2}{7} & \frac{3}{14} \\ \frac{1}{\sqrt{5}} & \frac{6}{\sqrt{70}} & -\frac{4}{7} & -\frac{3}{7} \\ 0 & \frac{5}{\sqrt{70}} & \frac{6}{7} & \frac{9}{14} \\ 0 & 0 & 1 & 2 \end{pmatrix}$, which now has its first two columns orthonormalized,

and orthogonal to the last two columns. We then normalize the third column by dividing

by $r_{33} = \sqrt{\frac{15}{7}}$, and so
$$\begin{pmatrix} \frac{2}{\sqrt{5}} & -\frac{3}{\sqrt{70}} & \frac{2}{\sqrt{105}} & \frac{3}{14} \\ \frac{1}{\sqrt{5}} & \frac{6}{\sqrt{70}} & -\frac{4}{\sqrt{105}} & -\frac{3}{7} \\ 0 & \frac{5}{\sqrt{70}} & \frac{6}{\sqrt{105}} & \frac{9}{14} \\ 0 & 0 & \frac{7}{\sqrt{105}} & 2 \end{pmatrix}.$$
 Finally, we subtract $r_{34} = \frac{20}{\sqrt{105}}$ times the third column from the fourth column. Dividing the resulting fourth column by its norm $r_{44} = \sqrt{\frac{5}{6}}$ results in the final formulas,

$$Q = \begin{pmatrix} \frac{2}{\sqrt{5}} & -\frac{3}{\sqrt{70}} & \frac{2}{\sqrt{105}} & -\frac{1}{\sqrt{30}} \\ \frac{1}{\sqrt{5}} & \frac{6}{\sqrt{70}} & -\frac{4}{\sqrt{105}} & \frac{2}{\sqrt{30}} \\ 0 & \frac{5}{\sqrt{70}} & \frac{6}{\sqrt{105}} & -\frac{3}{\sqrt{30}} \\ 0 & 0 & \frac{7}{\sqrt{105}} & \frac{4}{\sqrt{30}} \end{pmatrix}, \quad R = \begin{pmatrix} \sqrt{5} & \frac{4}{\sqrt{5}} & \frac{1}{\sqrt{5}} & 0 \\ 0 & \frac{\sqrt{14}}{\sqrt{5}} & \frac{16}{\sqrt{70}} & \frac{5}{\sqrt{70}} \\ 0 & 0 & \frac{\sqrt{15}}{\sqrt{7}} & \frac{20}{\sqrt{105}} \\ 0 & 0 & 0 & \frac{\sqrt{5}}{\sqrt{6}} \end{pmatrix},$$

for the $A = QR$ factorization.

6. Numerical Calculation of Eigenvalues.

We are now ready to apply these results to the problem of numerically approximating eigenvalues of matrices. Before explaining the QR algorithm, we first review the elementary *power method*.

The Power Method

We assume, for simplicity, that A is a *complete* $n \times n$ matrix, meaning that it has an eigenvector basis[†] $\mathbf{v}_1, \dots, \mathbf{v}_n$, with corresponding eigenvalues $\lambda_1, \dots, \lambda_n$. It is easy to see that the solution to the linear iterative system

$$\mathbf{v}^{(k+1)} = A\mathbf{v}^{(k)}, \quad \mathbf{v}^{(0)} = \mathbf{v}, \quad (6.1)$$

is obtained by multiplying the initial vector \mathbf{v} by the successive powers of the coefficient matrix: $\mathbf{v}^{(k)} = A^k \mathbf{v}$. If we write the initial vector in terms of the eigenvector basis

$$\mathbf{v} = c_1 \mathbf{v}_1 + \dots + c_n \mathbf{v}_n, \quad (6.2)$$

then the solution takes the explicit form

$$\mathbf{v}^{(k)} = A^k \mathbf{v} = c_1 \lambda_1^k \mathbf{v}_1 + \dots + c_n \lambda_n^k \mathbf{v}_n. \quad (6.3)$$

Suppose further that A has a single dominant *real* eigenvalue, λ_1 , that is larger than all others in magnitude, so

$$|\lambda_1| > |\lambda_j| \quad \text{for all } j > 1. \quad (6.4)$$

[†] This is not a very severe restriction. Theorem 4.1 implies that all symmetric matrices are complete. Moreover, perturbations caused by round off and/or numerical inaccuracies will almost inevitably make an incomplete matrix complete.

As its name implies, this eigenvalue will eventually dominate the iteration (6.3). Indeed, since

$$|\lambda_1|^k \gg |\lambda_j|^k \quad \text{for all } j > 1 \quad \text{and all } k \gg 0,$$

the first term in the iterative formula (6.3) will eventually be much larger than the rest, and so, provided $c_1 \neq 0$,

$$\mathbf{v}^{(k)} \approx c_1 \lambda_1^k \mathbf{v}_1 \quad \text{for } k \gg 0.$$

Therefore, the solution to the iterative system (6.1) will, almost always, end up being a multiple of the dominant eigenvector of the coefficient matrix.

To compute the corresponding eigenvalue, we note that the i^{th} entry of the iterate $\mathbf{v}^{(k)}$ is approximated by $v_i^{(k)} \approx c_1 \lambda_1^k v_{1,i}$, where $v_{1,i}$ is the i^{th} entry of the eigenvector \mathbf{v}_1 . Thus, as long as $v_{1,i} \neq 0$, we can recover the dominant eigenvalue by taking a ratio between selected components of successive iterates:

$$\lambda_1 \approx \frac{v_i^{(k)}}{v_i^{(k-1)}}, \quad \text{provided that } v_i^{(k-1)} \neq 0. \quad (6.5)$$

Example 6.1. Consider the matrix $A = \begin{pmatrix} -1 & 2 & 2 \\ -1 & -4 & -2 \\ -3 & 9 & 7 \end{pmatrix}$. As you can check, its eigenvalues and eigenvectors are

$$\lambda_1 = 3, \quad \mathbf{v}_1 = \begin{pmatrix} 1 \\ -1 \\ 3 \end{pmatrix}, \quad \lambda_2 = -2, \quad \mathbf{v}_2 = \begin{pmatrix} 0 \\ 1 \\ -1 \end{pmatrix}, \quad \lambda_3 = 1, \quad \mathbf{v}_3 = \begin{pmatrix} -1 \\ 1 \\ -2 \end{pmatrix}.$$

Repeatedly multiplying an initial vector $\mathbf{v} = (1, 0, 0)^T$, say, by A results in the iterates $\mathbf{v}^{(k)} = A^k \mathbf{v}$ listed in the accompanying table. The last column indicates the ratio $\lambda^{(k)} = v_1^{(k)} / v_1^{(k-1)}$ between the first components of successive iterates. (One could equally well use the second or third components.) The ratios are converging to the dominant eigenvalue $\lambda_1 = 3$, while the vectors $\mathbf{v}^{(k)}$ are converging to a very large multiple of the corresponding eigenvector $\mathbf{v}_1 = (1, -1, 3)^T$.

Variants of the power method for finding other eigenvalues include the *inverse power method* based on iterating the inverse matrix A^{-1} , with eigenvalues $\lambda_1^{-1}, \dots, \lambda_n^{-1}$ and the *shifted inverse power method*, based on $(A - \mu I)^{-1}$, with eigenvalues $(\lambda_k - \mu)^{-1}$. The power method only produces the dominant (largest in magnitude) eigenvalue of a matrix A . The inverse power method can be used to find the smallest eigenvalue. Additional eigenvalues can be found by using the shifted inverse power method, or deflation. However, if we need to know *all* the eigenvalues, such piecemeal methods are too time-consuming to be of much practical value.

The QR Algorithm

The most popular scheme for simultaneously approximating all the eigenvalues of a matrix A is the remarkable QR algorithm, first proposed in 1961 independently by Francis

and Kublanovskaya. The underlying idea is simple, but surprising. The first step is to factor the matrix

$$A = A_1 = Q_1 R_1$$

into a product of an orthogonal matrix Q_1 and a positive (i.e., with all positive entries along the diagonal) upper triangular matrix R_1 . Next, multiply the two factors together *in the wrong order!* The result is the new matrix

$$A_2 = R_1 Q_1.$$

We then repeat these two steps. Thus, we next factor

$$A_2 = Q_2 R_2$$

using the Gram–Schmidt process, and then multiply the factors in the reverse order to produce

$$A_3 = R_2 Q_2.$$

The complete algorithm can be written as

$$A = Q_1 R_1, \quad R_k Q_k = A_{k+1} = Q_{k+1} R_{k+1}, \quad k = 1, 2, 3, \dots, \quad (6.6)$$

where Q_k, R_k come from the previous step, and the subsequent orthogonal matrix Q_{k+1} and positive upper triangular matrix R_{k+1} are computed by using the numerically stable form of the Gram–Schmidt algorithm.

The astonishing fact is that, for many matrices A , the iterates $A_k \rightarrow V$ converge to an upper triangular matrix V whose diagonal entries are the eigenvalues of A . Thus, after a sufficient number of iterations, say k^* , the matrix A_{k^*} will have very small entries below the diagonal, and one can read off a complete system of (approximate) eigenvalues along its diagonal. For each eigenvalue, the computation of the corresponding eigenvector can be done by solving the appropriate homogeneous linear system, or by applying the shifted inverse power method.

Example 6.2. Consider the matrix $A = \begin{pmatrix} 2 & 1 \\ 2 & 3 \end{pmatrix}$. The initial Gram–Schmidt factorization $A = Q_1 R_1$ yields

$$Q_1 = \begin{pmatrix} .7071 & -.7071 \\ .7071 & .7071 \end{pmatrix}, \quad R_1 = \begin{pmatrix} 2.8284 & 2.8284 \\ 0 & 1.4142 \end{pmatrix}.$$

These are multiplied in the reverse order to give

$$A_2 = R_1 Q_1 = \begin{pmatrix} 4 & 0 \\ 1 & 1 \end{pmatrix}.$$

We refactor $A_2 = Q_2 R_2$ via Gram–Schmidt, and then reverse multiply to produce

$$Q_2 = \begin{pmatrix} .9701 & -.2425 \\ .2425 & .9701 \end{pmatrix}, \quad R_2 = \begin{pmatrix} 4.1231 & .2425 \\ 0 & .9701 \end{pmatrix},$$

$$A_3 = R_2 Q_2 = \begin{pmatrix} 4.0588 & -.7647 \\ .2353 & .9412 \end{pmatrix}.$$

The next iteration yields

$$Q_3 = \begin{pmatrix} .9983 & -.0579 \\ .0579 & .9983 \end{pmatrix}, \quad R_3 = \begin{pmatrix} 4.0656 & -.7090 \\ 0 & .9839 \end{pmatrix},$$

$$A_4 = R_3 Q_3 = \begin{pmatrix} 4.0178 & -.9431 \\ .0569 & .9822 \end{pmatrix}.$$

Continuing in this manner, after 9 iterations we find, to four decimal places,

$$Q_{10} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, \quad R_{10} = \begin{pmatrix} 4 & -1 \\ 0 & 1 \end{pmatrix}, \quad A_{11} = R_{10} Q_{10} = \begin{pmatrix} 4 & -1 \\ 0 & 1 \end{pmatrix}.$$

The eigenvalues of A , namely 4 and 1, appear along the diagonal of A_{11} . Additional iterations produce very little further change, although they can be used for increasing the accuracy of the computed eigenvalues.

If the original matrix A is symmetric, positive definite, and with distinct eigenvalues, then, in most situations (see below for the precise requirement), the iterates converge, $A_k \rightarrow V = \Lambda$, to a diagonal matrix containing the eigenvalues of A in decreasing order. Moreover, if, in this case, we recursively define

$$S_1 = Q_1, \quad S_k = S_{k-1} Q_k = Q_1 Q_2 \cdots Q_{k-1} Q_k, \quad k > 1, \quad (6.7)$$

then $S_k \rightarrow S$ have, as their limit, an orthogonal matrix whose columns are the orthonormal eigenvector basis of A .

Example 6.3. Consider the symmetric matrix $A = \begin{pmatrix} 2 & 1 & 0 \\ 1 & 3 & -1 \\ 0 & -1 & 6 \end{pmatrix}$. The initial $A = Q_1 R_1$ factorization produces

$$S_1 = Q_1 = \begin{pmatrix} .8944 & -.4082 & -.1826 \\ .4472 & .8165 & .3651 \\ 0 & -.4082 & .9129 \end{pmatrix}, \quad R_1 = \begin{pmatrix} 2.2361 & 2.2361 & -.4472 \\ 0 & 2.4495 & -3.2660 \\ 0 & 0 & 5.1121 \end{pmatrix},$$

and so

$$A_2 = R_1 Q_1 = \begin{pmatrix} 3.0000 & 1.0954 & 0 \\ 1.0954 & 3.3333 & -2.0870 \\ 0 & -2.0870 & 4.6667 \end{pmatrix}.$$

We refactor $A_2 = Q_2 R_2$ and reverse multiply to produce

$$Q_2 = \begin{pmatrix} .9393 & -.2734 & -.2071 \\ .3430 & .7488 & .5672 \\ 0 & -.6038 & .7972 \end{pmatrix}, \quad S_2 = S_1 Q_2 = \begin{pmatrix} .7001 & -.4400 & -.5623 \\ .7001 & .2686 & .6615 \\ -.1400 & -.8569 & .4962 \end{pmatrix},$$

$$R_2 = \begin{pmatrix} 3.1937 & 2.1723 & -.7158 \\ 0 & 3.4565 & -4.3804 \\ 0 & 0 & 2.5364 \end{pmatrix}, \quad A_3 = R_2 Q_2 = \begin{pmatrix} 3.7451 & 1.1856 & 0 \\ 1.1856 & 5.2330 & -1.5314 \\ 0 & -1.5314 & 2.0219 \end{pmatrix}.$$

Continuing in this manner, after 10 iterations we find

$$Q_{11} = \begin{pmatrix} 1.0000 & -.0067 & 0 \\ .0067 & 1.0000 & .0001 \\ 0 & -.0001 & 1.0000 \end{pmatrix}, \quad S_{11} = \begin{pmatrix} .0753 & -.5667 & -.8205 \\ .3128 & -.7679 & .5591 \\ -.9468 & -.2987 & .1194 \end{pmatrix},$$

$$R_{11} = \begin{pmatrix} 6.3229 & .0647 & 0 \\ 0 & 3.3582 & -.0006 \\ 0 & 0 & 1.3187 \end{pmatrix}, \quad A_{12} = \begin{pmatrix} 6.3232 & .0224 & 0 \\ .0224 & 3.3581 & -.0002 \\ 0 & -.0002 & 1.3187 \end{pmatrix}.$$

After 20 iterations, the process has completely settled down, and

$$Q_{21} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \quad S_{21} = \begin{pmatrix} .0710 & -.5672 & -.8205 \\ .3069 & -.7702 & .5590 \\ -.9491 & -.2915 & .1194 \end{pmatrix},$$

$$R_{21} = \begin{pmatrix} 6.3234 & .0001 & 0 \\ 0 & 3.3579 & 0 \\ 0 & 0 & 1.3187 \end{pmatrix}, \quad A_{22} = \begin{pmatrix} 6.3234 & 0 & 0 \\ 0 & 3.3579 & 0 \\ 0 & 0 & 1.3187 \end{pmatrix}.$$

The eigenvalues of A appear along the diagonal of A_{22} , while the columns of S_{21} are the corresponding orthonormal eigenvector basis, listed in the same order as the eigenvalues, both correct to 4 decimal places.

We will devote the remainder of this section to a justification of the QR algorithm for a general class of symmetric matrices. We begin by assuming that A is an $n \times n$ symmetric, positive definite matrix, with distinct eigenvalues

$$\lambda_1 > \lambda_2 > \cdots > \lambda_n > 0. \quad (6.8)$$

According to the Spectral Theorem 4.3, the corresponding unit eigenvectors $\mathbf{u}_1, \dots, \mathbf{u}_n$ form an orthonormal basis of \mathbb{R}^n . After treating this case, the analysis will be extended to a broader class of symmetric matrices, having both positive and negative eigenvalues.

The secret is that the QR algorithm is, in fact, a well-disguised adaptation of the more primitive power method. If we were to use the power method to capture all the eigenvectors and eigenvalues of A , the first thought might be to try to perform it simultaneously on a complete basis $\mathbf{v}_1^{(0)}, \dots, \mathbf{v}_n^{(0)}$ of \mathbb{R}^n instead of just one individual vector. The problem is that, for almost all vectors, the power iterates $\mathbf{v}_j^{(k)} = A^k \mathbf{v}_j^{(0)}$ all tend to a multiple of the dominant eigenvector \mathbf{u}_1 . Normalizing the vectors at each step is not any better, since then they merely converge to one of the two dominant unit eigenvectors $\pm \mathbf{u}_1$. However, if, inspired by the form of the eigenvector basis, we *orthonormalize* the vectors at each step, then we effectively prevent them from all accumulating at the same dominant unit eigenvector, and so, with some luck, the resulting vectors will converge to the full system of eigenvectors. Since orthonormalizing a basis via the Gram–Schmidt process is equivalent to a QR matrix factorization, the mechanics of the algorithm is not so surprising.

In detail, we start with any orthonormal basis, which, for simplicity, we take to be the standard basis vectors of \mathbb{R}^n , and so $\mathbf{u}_1^{(0)} = \mathbf{e}_1, \dots, \mathbf{u}_n^{(0)} = \mathbf{e}_n$. At the k^{th} stage of the algorithm, we set $\mathbf{u}_1^{(k)}, \dots, \mathbf{u}_n^{(k)}$ to be the orthonormal vectors that result from

applying the Gram–Schmidt algorithm to the power vectors $\mathbf{v}_j^{(k)} = A^k \mathbf{e}_j$. In matrix language, the vectors $\mathbf{v}_1^{(k)}, \dots, \mathbf{v}_n^{(k)}$ are merely the columns of A^k , and the orthonormal basis $\mathbf{u}_1^{(k)}, \dots, \mathbf{u}_n^{(k)}$ are the columns of the orthogonal matrix S_k in the QR decomposition of the k^{th} power of A , which we denote by

$$A^k = S_k P_k, \quad (6.9)$$

where P_k is positive upper triangular. Note that, in view of (6.6)

$$\begin{aligned} A &= Q_1 R_1, & A^2 &= Q_1 R_1 Q_1 R_1 = Q_1 Q_2 R_2 R_1, \\ A^3 &= Q_1 R_1 Q_1 R_1 Q_1 R_1 = Q_1 Q_2 R_2 Q_2 R_2 R_1 = Q_1 Q_2 Q_3 R_3 R_2 R_1, \end{aligned}$$

and, in general,

$$A^k = (Q_1 Q_2 \cdots Q_{k-1} Q_k) (R_k R_{k-1} \cdots R_2 R_1). \quad (6.10)$$

The product of orthogonal matrices is also orthogonal. The product of positive upper triangular matrices is also positive upper triangular. Therefore, comparing (6.9, 10) and invoking the uniqueness of the QR factorization, we conclude that

$$S_k = Q_1 Q_2 \cdots Q_{k-1} Q_k = S_{k-1} Q_k, \quad P_k = R_k R_{k-1} \cdots R_2 R_1 = R_k P_{k-1}. \quad (6.11)$$

Let $S = (\mathbf{u}_1 \ \mathbf{u}_2 \ \dots \ \mathbf{u}_n)$ denote an orthogonal matrix whose columns are unit eigenvectors of A . The Spectral Theorem 4.3 tells us that

$$A = S \Lambda S^T, \quad \text{where} \quad \Lambda = \text{diag}(\lambda_1, \dots, \lambda_n)$$

is the diagonal eigenvalue matrix. Substituting the spectral factorization into (6.9), we find

$$A^k = S \Lambda^k S^T = S_k P_k.$$

We now make one additional assumption on the matrix A by requiring that S^T be a *regular matrix*, meaning that we can factor $S^T = LU$ into a product of special lower and upper triangular matrices, or, equivalently, that Gaussian Elimination can be performed on the linear system $S^T \mathbf{x} = \mathbf{b}$ *without any row interchanges*. Regularity holds generically, and is the analog of the condition that our initial vector in the power method includes a nonzero component of the dominant eigenvector. We can assume that, without loss of generality, the diagonal entries of the upper triangular matrix U — that is, the pivots of S^T — are all positive. Indeed, this can be arranged by multiplying each row of S^T by the sign of its pivot — this amounts to possibly changing the signs of some of the unit eigenvectors \mathbf{u}_i , which is allowed since it does not affect their status as an orthonormal eigenvector basis.

Under these two assumptions,

$$A^k = S \Lambda^k LU = S_k P_k, \quad \text{and hence} \quad S \Lambda^k L = S_k P_k U^{-1}.$$

Multiplying on the right by Λ^{-k} , we obtain

$$S \Lambda^k L \Lambda^{-k} = S_k T_k, \quad \text{where} \quad T_k = P_k U^{-1} \Lambda^{-k} \quad (6.12)$$

is also a positive upper triangular matrix, since we are assuming that the eigenvalues, i.e., the diagonal entries of Λ , are all positive.

Now consider what happens as $k \rightarrow \infty$. The entries of the lower triangular matrix $N = \Lambda^k L \Lambda^{-k}$ are

$$n_{ij} = \begin{cases} l_{ij}(\lambda_i/\lambda_j)^k, & i > j, \\ l_{ii} = 1, & i = j, \\ 0, & i < j. \end{cases}$$

Since we are assuming $\lambda_i < \lambda_j$ when $i > j$, we immediately deduce that

$$\Lambda^k L \Lambda^{-k} \longrightarrow \mathbf{I},$$

and hence

$$S_k T_k = S \Lambda^k L \Lambda^{-k} \longrightarrow S \quad \text{as} \quad k \longrightarrow \infty. \quad (6.13)$$

We now appeal to the following lemma, whose proof will be given after we finish the justification of the QR algorithm.

Lemma 6.4. *Let S_1, S_2, \dots and S be orthogonal matrices and T_1, T_2, \dots positive upper triangular matrices. Then $S_k T_k \rightarrow S$ as $k \rightarrow \infty$ if and only if $S_k \rightarrow S$ and $T_k \rightarrow \mathbf{I}$.*

Therefore, as claimed, the orthogonal matrices S_k do converge to the orthogonal eigenvector matrix S . By (6.11),

$$Q_k = S_{k-1}^{-1} S_k \longrightarrow \mathbf{I} \quad \text{as} \quad k \longrightarrow \infty. \quad (6.14)$$

Moreover, by (6.11–12),

$$R_k = P_k P_{k-1}^{-1} = (T_k \Lambda^k U^{-1}) (T_{k-1} \Lambda^{k-1} U^{-1})^{-1} = T_k \Lambda T_{k-1}^{-1}.$$

Since both T_k and T_{k-1} converge to the identity matrix, in the limit

$$R_k \longrightarrow \Lambda \quad \text{as} \quad k \longrightarrow \infty. \quad (6.15)$$

Combining (6.14, 15), we deduce that the matrices appearing in the QR algorithm converge to the diagonal eigenvalue matrix, as claimed:

$$A_k = Q_k R_k \longrightarrow \Lambda \quad \text{as} \quad k \longrightarrow \infty. \quad (6.16)$$

Note that the eigenvalues appear in decreasing order along the diagonal, which is a consequence of our regularity assumption on the transposed eigenvector matrix S^T .

The last remaining item is a proof of Lemma 6.4. We write $S = (\mathbf{u}_1 \ \mathbf{u}_2 \ \dots \ \mathbf{u}_n)$, $S_k = (\mathbf{u}_1^{(k)}, \dots, \mathbf{u}_n^{(k)})$ in columnar form. Let $t_{ij}^{(k)}$ denote the entries of the positive upper triangular matrix T_k . The first column of the limiting equation $S_k T_k \rightarrow S$ reads

$$t_{11}^{(k)} \mathbf{u}_1^{(k)} \longrightarrow \mathbf{u}_1.$$

Since both $\mathbf{u}_1^{(k)}$ and \mathbf{u}_1 are unit vectors, and $t_{11}^{(k)} > 0$,

$$\|t_{11}^{(k)} \mathbf{u}_1^{(k)}\| = t_{11}^{(k)} \longrightarrow \|\mathbf{u}_1\| = 1, \quad \text{and hence} \quad \mathbf{u}_1^{(k)} \longrightarrow \mathbf{u}_1.$$

The second column reads

$$t_{12}^{(k)} \mathbf{u}_1^{(k)} + t_{22}^{(k)} \mathbf{u}_2^{(k)} \longrightarrow \mathbf{u}_2.$$

Taking the inner product with $\mathbf{u}_1^{(k)} \rightarrow \mathbf{u}_1$ and using orthonormality, we deduce $t_{12}^{(k)} \rightarrow 0$, and so $t_{22}^{(k)} \mathbf{u}_2^{(k)} \rightarrow \mathbf{u}_2$, which, by the previous reasoning, implies $t_{22}^{(k)} \rightarrow 1$ and $\mathbf{u}_2^{(k)} \rightarrow \mathbf{u}_2$. The proof is completed by working in order through the remaining columns, employing a similar argument at each step.

To modify the preceding proof for symmetric, but non-positive definite matrices, we introduce the diagonal orthogonal matrix

$$\Delta = \Delta^{-1} = \text{diag}(\text{sign } \lambda_1, \dots, \text{sign } \lambda_n), \quad \text{so that} \quad \Delta \Lambda = \Lambda \Delta = \text{diag}(|\lambda_1|, \dots, |\lambda_n|).$$

whose diagonal entries are the signs of the eigenvalues. (So in the positive definite case $\Delta = \mathbf{I}$.) Set

$$\tilde{S}_k = S_k \Delta^k, \quad \tilde{T}_k = \Delta^k T_k.$$

Thus, (6.13) can be rewritten as

$$\tilde{S}_k \tilde{T}_k = S_k T_k = S \Lambda^k L \Lambda^{-k} \longrightarrow S \quad \text{as} \quad k \longrightarrow \infty.$$

Recalling (6.12), the diagonal entries of the upper triangular matrix $\tilde{T}_k = \Delta^k T_k$ are seen to be strictly positive, and hence Lemma 6.4 implies

$$\tilde{S}_k \longrightarrow S, \quad \tilde{T}_k \longrightarrow \mathbf{I}, \quad \text{as} \quad k \longrightarrow \infty.$$

Therefore,

$$\begin{aligned} Q_k &= S_{k-1}^{-1} S_k = \Delta^{k-1} \tilde{S}_{k-1}^{-1} \tilde{S}_k \Delta^k \longrightarrow \Delta, \\ R_k &= T_k \Lambda T_{k-1}^{-1} = \Delta^k \tilde{T}_k \Lambda \tilde{T}_{k-1}^{-1} \Delta^{k-1} \longrightarrow \Delta \Lambda, \end{aligned} \quad \text{as} \quad k \longrightarrow \infty. \quad (6.17)$$

The convergence of the QR matrices immediately follows:

$$A_k = Q_k R_k \longrightarrow \Delta^2 \Lambda = \Lambda \quad \text{as} \quad k \longrightarrow \infty.$$

Keep in mind that, in non-positive definite cases, the orthogonal matrices S_k no longer converge, since the nonzero entries in the columns corresponding to negative eigenvalues switch their signs back and forth as k increases; however, by eliminating the sign changes — which, in practice can be done by inspection — their modifications $\tilde{S}_k = S_k \Delta^k \rightarrow S$ will converge to the eigenvector matrix.

This completes the proof of the convergence of the QR algorithm for a broad class of symmetric matrices.

Theorem 6.5. *If A is symmetric, satisfies*

$$|\lambda_1| > |\lambda_2| > \dots > |\lambda_n| > 0, \quad (6.18)$$

and its transposed eigenvector matrix S^T is regular, then the matrices appearing in the QR algorithm applied to A converge to the diagonal eigenvalue matrix: $A_k \rightarrow \Lambda = \text{diag}(\lambda_1, \dots, \lambda_n)$ as $k \rightarrow \infty$.

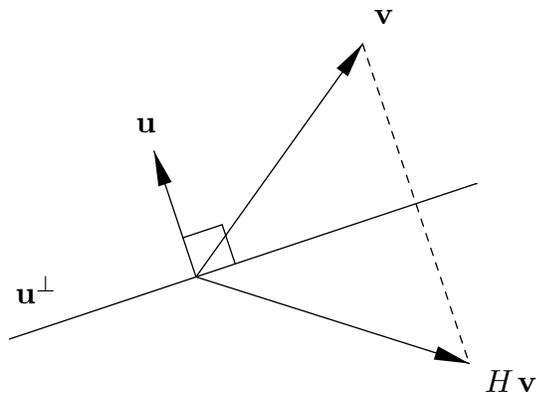


Figure 2. Elementary Reflection Matrix.

Tridiagonalization

In practical implementations, the direct QR algorithm often takes too long to provide reasonable approximations to the eigenvalues of large matrices. Fortunately, the algorithm can be made much more efficient by a simple preprocessing step. The key observation is that the QR algorithm preserves the class of symmetric tridiagonal matrices, and, moreover, like Gaussian Elimination, is much faster when applied to this class of matrices. It turns out that, although diagonalizing a matrix is effectively the same as finding its eigenvectors, one can “tridiagonalize” the matrix by a sequence of fairly elementary matrix operations, based on the following class of matrices.

Consider the *Householder* or *elementary reflection matrix*

$$H = I - 2\mathbf{u}\mathbf{u}^T \quad (6.19)$$

in which \mathbf{u} is a unit vector (in the Euclidean norm). The matrix H represents a reflection of vectors through the subspace $\mathbf{u}^\perp = \{\mathbf{v} \mid \mathbf{v} \cdot \mathbf{u} = 0\}$ of vectors orthogonal to \mathbf{u} , as illustrated in Figure 2. The matrix H is symmetric and orthogonal:

$$H^T = H, \quad H^2 = I, \quad H^{-1} = H. \quad (6.20)$$

The proof is straightforward: symmetry is immediate, while

$$H H^T = H^2 = (I - 2\mathbf{u}\mathbf{u}^T)(I - 2\mathbf{u}\mathbf{u}^T) = I - 4\mathbf{u}\mathbf{u}^T + 4\mathbf{u}(\mathbf{u}^T\mathbf{u})\mathbf{u}^T = I$$

since, by assumption, $\mathbf{u}^T\mathbf{u} = \|\mathbf{u}\|^2 = 1$. By suitably forming the unit vector \mathbf{u} , we can construct an elementary reflection matrix that interchanges any two vectors of the same length.

Lemma 6.6. *Let $\mathbf{v}, \mathbf{w} \in \mathbb{R}^n$ with $\|\mathbf{v}\| = \|\mathbf{w}\|$. Set $\mathbf{u} = (\mathbf{v} - \mathbf{w})/\|\mathbf{v} - \mathbf{w}\|$ and let $H = I - 2\mathbf{u}\mathbf{u}^T$ be the corresponding elementary reflection matrix. Then $H\mathbf{v} = \mathbf{w}$ and $H\mathbf{w} = \mathbf{v}$.*

Proof: Keeping in mind that \mathbf{x} and \mathbf{y} have the same Euclidean norm, we compute

$$\begin{aligned} H\mathbf{v} &= (\mathbf{I} - 2\mathbf{u}\mathbf{u}^T)\mathbf{v} = \mathbf{v} - 2\frac{(\mathbf{v} - \mathbf{w})(\mathbf{v} - \mathbf{w})^T\mathbf{v}}{\|\mathbf{v} - \mathbf{w}\|^2} \\ &= \mathbf{v} - 2\frac{(\mathbf{v} - \mathbf{w})(\|\mathbf{v}\|^2 - \mathbf{w} \cdot \mathbf{v})}{2\|\mathbf{v}\|^2 - 2\mathbf{v} \cdot \mathbf{w}} = \mathbf{v} - (\mathbf{v} - \mathbf{w}) = \mathbf{w}. \end{aligned}$$

The proof of the second equation is similar.

Q.E.D.

In Householder's approach to the QR factorization, we were able to convert the matrix A to upper triangular form R by a sequence of elementary reflection matrices. Unfortunately, this procedure does not preserve the eigenvalues of the matrix — the diagonal entries of R are *not* the eigenvalues — and so we need to be a bit more clever here. We begin by recalling that similar matrices have the same eigenvalues.

Lemma 6.7. *If $H = \mathbf{I} - 2\mathbf{u}\mathbf{u}^T$ is an elementary reflection matrix, with \mathbf{u} a unit vector, then A and $B = HAH$ are similar matrices and hence have the same eigenvalues.*

Proof: According to (6.20), $H^{-1} = H$, and hence $B = H^{-1}AH$ is similar to A . If λ is any eigenvalue of A , so $A\mathbf{v} = \lambda\mathbf{v}$ for $\mathbf{v} \neq \mathbf{0}$, then $B\mathbf{w} = \lambda\mathbf{w}$ where $\mathbf{w} = H^{-1}\mathbf{v}$, which implies that λ remains an eigenvalue of B , and conversely. *Q.E.D.*

Given a symmetric $n \times n$ matrix A , our goal is to devise a similar tridiagonal matrix by applying a sequence of Householder reflections. We begin by setting

$$\mathbf{x}_1 = \begin{pmatrix} 0 \\ a_{21} \\ a_{31} \\ \vdots \\ a_{n1} \end{pmatrix}, \quad \mathbf{y}_1 = \begin{pmatrix} 0 \\ \pm r_1 \\ 0 \\ \vdots \\ 0 \end{pmatrix}, \quad \text{where} \quad r_1 = \|\mathbf{x}_1\| = \|\mathbf{y}_1\|,$$

so that \mathbf{x}_1 contains all the off-diagonal entries of the first column of A . Let

$$H_1 = \mathbf{I} - 2\mathbf{u}_1\mathbf{u}_1^T, \quad \text{where} \quad \mathbf{u}_1 = \frac{\mathbf{x}_1 - \mathbf{y}_1}{\|\mathbf{x}_1 - \mathbf{y}_1\|}$$

be the corresponding elementary reflection matrix that maps \mathbf{x}_1 to \mathbf{y}_1 . Either \pm sign in the formula for \mathbf{y}_1 works in the algorithm; a good choice is to set it to be the opposite of the sign of the entry a_{21} , which helps minimize the possible effects of round-off error when computing the unit vector \mathbf{u}_1 . By direct computation, based on Lemma 6.6 and the fact that the first entry of \mathbf{u}_1 is zero,

$$A_2 = H_1 A H_1 = \begin{pmatrix} a_{11} & r_1 & 0 & \dots & 0 \\ r_1 & \tilde{a}_{22} & \tilde{a}_{23} & \dots & \tilde{a}_{2n} \\ 0 & \tilde{a}_{32} & \tilde{a}_{33} & \dots & \tilde{a}_{3n} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & \tilde{a}_{n2} & \tilde{a}_{n3} & \dots & \tilde{a}_{nn} \end{pmatrix} \quad (6.21)$$

for certain \tilde{a}_{ij} ; the explicit formulae are not needed. Thus, by a single Householder transformation, we convert A into a similar matrix A_2 whose first row and column are in tridiagonal form. We repeat the process on the lower right $(n-1) \times (n-1)$ submatrix of A_2 . We set

$$\mathbf{x}_2 = \begin{pmatrix} 0 \\ 0 \\ \tilde{a}_{32} \\ \tilde{a}_{42} \\ \vdots \\ \tilde{a}_{n2} \end{pmatrix}, \quad \mathbf{y}_1 = \begin{pmatrix} 0 \\ 0 \\ \pm r_2 \\ 0 \\ \vdots \\ 0 \end{pmatrix}, \quad \text{where} \quad r_2 = \|\mathbf{x}_2\| = \|\mathbf{y}_2\|,$$

and the \pm sign is chosen to be the opposite of that of \tilde{a}_{32} . Setting

$$H_2 = \mathbf{I} - 2\mathbf{u}_2\mathbf{u}_2^T, \quad \text{where} \quad \mathbf{u}_2 = \frac{\mathbf{x}_2 - \mathbf{y}_2}{\|\mathbf{x}_2 - \mathbf{y}_2\|},$$

we construct the similar matrix

$$A_3 = H_2 A_2 H_2 = \begin{pmatrix} a_{11} & r_1 & 0 & 0 & \dots & 0 \\ r_1 & \tilde{a}_{22} & r_2 & 0 & \dots & 0 \\ 0 & r_2 & \hat{a}_{33} & \hat{a}_{34} & \dots & \hat{a}_{3n} \\ 0 & 0 & \hat{a}_{43} & \hat{a}_{44} & \dots & \hat{a}_{4n} \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \hat{a}_{n3} & \hat{a}_{n4} & \dots & \hat{a}_{nn} \end{pmatrix}.$$

whose first two rows and columns are now in tridiagonal form. The remaining steps in the algorithm should now be clear. Thus, the final result is a tridiagonal matrix $T = A_n$ that has the *same eigenvalues* as the original symmetric matrix A . Let us illustrate the method by an example.

Example 6.8. To tridiagonalize $A = \begin{pmatrix} 4 & 1 & -1 & 2 \\ 1 & 4 & 1 & -1 \\ -1 & 1 & 4 & 1 \\ 2 & -1 & 1 & 4 \end{pmatrix}$, we begin with its first

column. We set $\mathbf{x}_1 = \begin{pmatrix} 0 \\ 1 \\ -1 \\ 2 \end{pmatrix}$, so that $\mathbf{y}_1 = \begin{pmatrix} 0 \\ \sqrt{6} \\ 0 \\ 0 \end{pmatrix} \approx \begin{pmatrix} 0 \\ 2.4495 \\ 0 \\ 0 \end{pmatrix}$. Therefore, the unit

vector is $\mathbf{u}_1 = \frac{\mathbf{x}_1 - \mathbf{y}_1}{\|\mathbf{x}_1 - \mathbf{y}_1\|} = \begin{pmatrix} 0 \\ .8391 \\ -.2433 \\ .4865 \end{pmatrix}$, with corresponding Householder matrix

$$H_1 = \mathbf{I} - 2\mathbf{u}_1\mathbf{u}_1^T = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & -.4082 & .4082 & -.8165 \\ 0 & .4082 & .8816 & .2367 \\ 0 & -.8165 & .2367 & .5266 \end{pmatrix}.$$

Thus,

$$A_2 = H_1 A H_1 = \begin{pmatrix} 4.0000 & -2.4495 & 0 & 0 \\ -2.4495 & 2.3333 & -.3865 & -.8599 \\ 0 & -.3865 & 4.9440 & -.1246 \\ 0 & -.8599 & -.1246 & 4.7227 \end{pmatrix}.$$

In the next phase, $\mathbf{x}_2 = \begin{pmatrix} 0 \\ 0 \\ -.3865 \\ -.8599 \end{pmatrix}$, $\mathbf{y}_2 = \begin{pmatrix} 0 \\ 0 \\ -.9428 \\ 0 \end{pmatrix}$, so $\mathbf{u}_2 = \begin{pmatrix} 0 \\ 0 \\ -.8396 \\ -.5431 \end{pmatrix}$, and

$$H_2 = I - 2\mathbf{u}_2\mathbf{u}_2^T = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & -.4100 & -.9121 \\ 0 & 0 & -.9121 & .4100 \end{pmatrix}.$$

The resulting matrix

$$T = A_3 = H_2 A_2 H_2 = \begin{pmatrix} 4.0000 & -2.4495 & 0 & 0 \\ -2.4495 & 2.3333 & .9428 & 0 \\ 0 & .9428 & 4.6667 & 0 \\ 0 & 0 & 0 & 5 \end{pmatrix}$$

is now in tridiagonal form.

Since the final tridiagonal matrix T has the same eigenvalues as A , we can apply the QR algorithm to T to approximate the common eigenvalues. (The eigenvectors must then be computed separately, e.g., by the shifted inverse power method.) It is not hard to show that, if $A = A_1$ is tridiagonal, so are all the iterates A_2, A_3, \dots . Moreover, far fewer arithmetic operations are required. For instance, in the preceding example, after we apply 20 iterations of the QR algorithm directly to T , the upper triangular factor has become

$$R_{21} = \begin{pmatrix} 6.0000 & -.0065 & 0 & 0 \\ 0 & 4.5616 & 0 & 0 \\ 0 & 0 & 5.0000 & 0 \\ 0 & 0 & 0 & .4384 \end{pmatrix}.$$

The eigenvalues of T , and hence also of A , appear along the diagonal, and are correct to 4 decimal places.

Finally, even if A is not symmetric, one can still apply the same sequence of Householder transformations to simplify it. The final result is no longer tridiagonal, but rather a similar *upper Hessenberg matrix*, which means that all entries below the subdiagonal are zero, but those above the superdiagonal are not necessarily zero. For instance, a 5×5 upper Hessenberg matrix looks like

$$\begin{pmatrix} * & * & * & * & * \\ * & * & * & * & * \\ 0 & * & * & * & * \\ 0 & 0 & * & * & * \\ 0 & 0 & 0 & * & * \end{pmatrix},$$

where the starred entries can be anything. It can be proved that the QR algorithm maintains the upper Hessenberg form, and, while not as efficient as in the tridiagonal case, still yields a significant savings in computational effort required to find the common eigenvalues.

Acknowledgments: I would like to thank Raphaële Herbin for comments on an earlier version of these notes.